How cosmic rays mediate the evolution of the interstellar medium

Christine M. Simpson [1,2]★ Rüdiger Pakmor [3] Christoph Pfrommer [4] Simon C. O. Glover [5] and Rowan Smith [6]

[1] *Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA*
[2] *Department of Astronomy and Astrophysics, The University of Chicago, Chicago, IL 60637, USA*
[3] *Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany*
[4] *Leibniz-Institute for Astrophysics Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany*
[5] *Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik, Albert-Ueberle-Straße 2, D-69120 Heidelberg, Germany*
[6] *Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Oxford Road, Manchester M13 9PL, UK*

ABSTRACT

We explore the impact of diffusive cosmic rays (CRs) on the evolution of the interstellar medium (ISM) under varying assumptions of supernova explosion environment. In practice, we systematically vary the relative fractions of supernovae (SN) occurring in star-forming high-density gas and those occurring in random locations decoupled from star-forming gas to account for SN from run-away stars or explosions in regions that have been cleared by prior SN, stellar winds, or radiation. We find that in the simple system of a periodic stratified gas layer the ISM structure will evolve to one of two solutions: a 'peak driving' state where warm gas is volume filling or a 'thermal runaway' state where hot gas is volume filling. CR pressure and transport are important factors that strongly influence the solution state the ISM reaches and have the ability to flip the ISM between solutions. Observable signatures such as gamma-ray emission and H I gas are explored. We find that gamma-ray luminosity from pion decay is largely consistent with observations for a range of model parameters. The thickness of the H I gas layer may be too compact, however, this may be due to a large cold neutral fraction of mid-plane gas. The volume fraction of hot gas evolves to stable states in both solutions, but neither settles to a Milky Way-like configuration, suggesting that additional physics omitted here (e.g. a cosmological circumgalactic medium, radiation transport, or spectrally resolved and spatially varying CR transport) may be required.

Key words: MHD – cosmic rays – ISM: evolution – ISM: structure – gamma-rays: ISM.

1 INTRODUCTION

Observations and modelling of the interstellar medium (ISM) have long indicated that the structure of the ISM is heavily influenced by the explosion of massive stars as supernovae (SN) (e.g. McKee & Ostriker 1977). The phase structure of the ISM, where mass is aggregated in dense molecular clouds, embedded in a hot diffuse medium in pressure equilibrium, has been established through observations of neutral atomic and molecular gas (e.g. Cox 2005) and observations of diffuse ionized gas (e.g. Haffner et al. 2003; Long et al. 2022).

Modelling the phase structure and energetics of the ISM in a dynamic way has been a goal of simulations, both to understand the conditions for star formation within galaxies and to inform global models of galaxy evolution. Simulations on the ISM scale have demonstrated the importance of SN explosion environment (e.g. Iffrig & Hennebelle 2015; Walch et al. 2015; Ohlin, Renaud & Agertz 2019; Andersson, Agertz & Renaud 2020). Simple 1D models of SN shocks point to the importance of the background ISM gas density as a determining factor for how the SN shocks develop and the amount of energy and momentum they deposit in the ISM (Draine 2011). SN

explosions that occur in dense molecular clouds will have most of their energy radiatively cooled on short time-scales, but energy from explosions that occur in low density gas (due to pre-SN feedback, stellar kicks, or long SN delay times that allow for the dissipation of the natal molecular cloud) will remain in the ISM longer due to longer radiative cooling times (e.g. Walch & Naab 2015).

These simulations have most commonly studied the injection and dissipation of SN energy in a medium that only self-consistently tracks thermal and kinetic energies. However, the energy contained in the ISM comes in several other forms. Indeed, the magnetic fields threading through the gas also contain energy, as do cosmic rays (CRs) in the medium, which are high energy protons, electrons, and other heavy atomic nuclei and represent an additional reservoir of energy. Traditionally, these four energies are assumed to be in equilibrium (Boulares & Cox 1990; Cox 2005; Naab & Ostriker 2017), but this assumption would have to be scrutinized with numerical simulations of the interstellar medium.

The non-thermal energies from magnetic fields and CRs are of particular interest because they transport and dissipate energy differently from the thermal gas. CRs in particular have been posited as a mechanism for driving galactic outflows and as a regulator for star formation on interstellar scales (Girichidis et al. 2016, 2018; Simpson et al. 2016; Farber et al. 2018; Commerçon, Marcowith & Dubois 2019) as well as on galactic scales in isolated galaxies (Uhlig et al.

2012; Booth et al. 2013; Salem & Bryan 2014; Pakmor et al. 2016b; Ruszkowski, Yang & Zweibel 2017; Butsky & Quinn 2018; Jacob et al. 2018; Dashyan & Dubois 2020; Farcy et al. 2022; Girichidis et al. 2022; Quataert, Thompson & Jiang 2022) and in cosmological settings (Jubelgas et al. 2008; Salem, Bryan & Hummels 2014; Buck et al. 2020; Hopkins et al. 2021b). The physics by which galaxies retain or expel gas and the efficiency with which they then turn gas into stars is of cosmological importance (e.g. Naab & Ostriker 2017; Wechsler & Tinker 2018), but the feedback processes that drive these behaviours are still uncertain. SN explosions deposit only 5–15 per cent of their energy in CRs (Caprioli & Spitkovsky 2014; Pais et al. 2018), but CR energy dissipates on longer time-scales and couples to the thermal gas differently.

Substantial simulation work has been done to explore how CRs on the ISM scale impact the local launching of outflows, but less has been done to explore how it affects the structure of the ISM. Specifically, how CR pressure and transport impacts the picture of ISM structure driven by SN environment that has been built-up through the past decade. Rathjen et al. (2021) have studied how a myriad of feedback effects including stellar winds, radiation, and CRs impact ISM structure and the development of outflows. They find that CRs only have a moderate impact on the multiphase structure of the ISM, but their models are only evolved for 100 Myr and they do not explore details of CR pressure or differences in transport.

The impact of ISM structure on observables associated with CRs, such as gamma-ray emission or radio continuum emission, also have yet to be studied. CR observables on galactic scales in the Milky Way (Moskalenko & Strong 1998; Evoli et al. 2008; Kissmann et al. 2015; Hopkins et al. 2021a; Werhahn et al. 2021a) and non-thermal radio and gamma-ray emission from galaxies are a promising avenue for probing CR calorimetry (which fraction of CR energy is radiated away as opposed to providing dynamical feedback; Thompson et al. 2006; Lacki, Thompson & Quataert 2010; Pfrommer et al. 2021; Werhahn et al. 2021b; Werhahn, Pfrommer & Girichidis 2021c) and of constraining the physics of CR transport (Thomas, Pfrommer & Enßlin 2020).

We aim to address these issues and specifically to explore how the presence of CRs impacts the structure and evolution of the ISM driven by SN explosions. This paper is part of a series utilizing a stratified box setup to explore the impact of CRs on the ISM. We test different models for SN explosion placement in the presence of CRs and demonstrate how the presence of an energy reservoir with different transport and dissipation time-scales alters the ISM structure. We quantify how CR energy is dissipated from the galactic mid-plane and explore models with different diffusivity properties. We finally explore observable implications of our models in gamma-ray and 21 cm emission.

The paper is organized as follows: Section 2 describes our methodology and simulation set-up, including our model for SN explosion placement; Section 3.1 describes the effect SN explosion placement has on ISM structure when CRs are included; Sections 3.2 and 3.3 describe the energetics of mid-plane gas; Section 3.4 describes how our results depend on differences in the CR model; and Section 3.5 explores the impact of CRs on the mid-plane pressure. In Section 4 we look at observable signatures of these models in gamma-ray and 21 cm emission. Finally, we conclude with a discussion of the MW's ISM structure and how it compares to the models presented here.

2 METHODS

We use the second-order accurate moving-mesh code AREPO to study the impact of SN explosions and CR transport on the ISM (Springel

2010; Pakmor et al. 2016a). AREPO gives a quasi-Lagrangian solution to the ideal (magneto-)hydrodynamics equations that combines advantages of an Eulerian scheme in capturing shocks and discontinuities and the advantages of a pure Lagrangian scheme with improved Galilean invariance. Here we describe the set-up and initial conditions of our simulations; a model for pressure in high density gas; additional gas physics modules such as self-gravity, radiative cooling, cosmic rays, and MHD; and the different schemes for SN explosion placement that we explore.

2.1 Simulation setup and initial conditions

We simulate the impact of SN explosions on the ISM in a tall column of stratified gas meant to capture the behaviour of gas in a small patch of a galactic disc. The simulation domain has a width of $L = 1$ kpc in both the horizontal x and y directions and extends ± 5 kpc in the vertical z direction (for a box height of $L_z = 10$ kpc). The domain has periodic boundaries in the x and y directions and outflow boundaries in the z direction.

Mass and energy can flow out of the simulation volume through the outflow boundaries at the top and bottom of the column. This is accomplished by fixing a layer of static 'sticky' cells along these boundaries that do not move (i.e. their positions are not updated by the moving mesh algorithm). This is necessary to prevent the mesh structure (that normally moves with the motion of the thermal gas) from exiting the simulation volume. A layer of 'ghost' cells are mirrored across the outflow boundaries and gas can move across this boundary in the outward direction only, leaving the simulation.

The initial positions and densities of gas cells are set according to an assumption of isothermal hydrostatic equilibrium following Creasey, Theuns & Bower (2013) and used in Simpson et al. (2016), the preliminary study to this work. The initial density of gas cells depends on the choice of initial gas surface density, Σ_0, and gas fraction, f_g, as well as the cell's height h above the mid-plane:

$$\rho_0(h) = \frac{\Sigma_0}{2b_0} \text{sech}^2 \left(\frac{h}{b_0} \right), \tag{1}$$

$$b_0 = \frac{f_g k_B T_0}{m_p \mu_0 \pi G \Sigma_0}. \tag{2}$$

The constant G is the gravitational constant, k_B is Boltzmann's constant, and m_p is the proton mass. The variable T_0 is the initial isothermal temperature, chosen to be 10^4 K, and μ_0 is the initial mean molecular weight, chosen to be 0.61. A lower limit to the initial gas density of 10^{-20} $M_\odot$ pc^{-3} is imposed affecting cells at high h. We choose values of $\Sigma_0 = 10$ $M_\odot$ pc^{-2} and $f_g = 0.1$, which are similar to the gas surface density and gas fraction of the Milky Way's disc at the solar circle (Flynn et al. 2006). The quantity b_0 is the isothermal scale height of the disc and is 100 pc for the models presented here. The height containing half the mass in the box is related to b_0 by $h_{1/2} = 0.55b_0$.

Gas cells are part of a Voronoi mesh that is constructed around a set of 'mesh-generating' points. The bulk of the mesh-generating points are distributed with uniform randomness in the horizontal x-y plane and drawn randomly from a distribution function that follows equation (1) in the z direction. The number of points placed in this fashion is $\Sigma_0 L^2 m_t^{-1}$, where m_t is the target gas mass that the mesh refinement scheme will seek to maintain during the simulation. For our fiducial model with $\Sigma_0 = 10$ $M_\odot$ pc^{-2}, $L = 1$ kpc, and $m_t = 10$ $M_\odot$, this number is 10^6 points. These points are supplemented by a Cartesian background mesh that fills the entire volume and has a cell diameter of 50 pc. The Cartesian mesh provides an initial minimum resolution for the gas in low-density regions at high h.

Once the simulation begins, the Cartesian mesh deforms according to the Voronoi mesh generation scheme of AREPO.

Once these mesh-generating points are placed, and before our science simulations are run, we run the initial conditions with AREPO in the 'mesh-relax' mode. This mode does not solve gravity, hydrodynamics, or any other physics module, but rather, the code only evolves the Voronoi mesh using the mesh-regularization algorithm (Lloyd 1982; Springel 2010). This algorithm seeks to make cells more 'round' by moving mesh generating points closer to their cells' centres of mass. We typically run the initial mesh for approximately 100 time-steps in this mode with a slightly reduced maximum cell opening angle factor of 1.75 (in our science runs, we return to the default AREPO value of 2.25) until the mesh no longer changes. After, gas densities are reassigned to the new relaxed mesh based on cells' new height h. This step eliminates awkwardly shaped cells that can result from the initial random placement of mesh generating points. It has the effect of regularizing the mesh in the mid-plane, but leaves the Cartesian mesh at large h largely unchanged.

2.2 Hydrodynamics and magnetohydrodynamics

To solve the ideal hydrodynamics equations, we use the second-order accurate moving-mesh code AREPO (Springel 2010; Pakmor et al. 2016a) with a thermal adiabatic index γ of 5/3.[1] We have a target gas mass for cells of $10\,M_\odot$ and we impose a maximum time-step of 0.1 Myr. In many of our simulations, ideal magnetohydrodynamics (MHD) is included with the Powell 8-wave scheme for divergence control (Pakmor & Springel 2013). When using MHD, an initial seed field that varies with height h above the mid-plane is adopted with a strength of $10^{-7}\,G \times \text{sech}^{4/3}(h/b_0)$, which scales the field as $\rho_0^{2/3}$ following equation (1). Its initial orientation is purely horizontal and parallel to the gas mid-plane, pointing in the x direction. This initial configuration is divergence free.

2.2.1 Cooling and chemistry

Radiative cooling and chemistry are modelled following Smith et al. (2014) who use atomic and molecular hydrogen chemistry and cooling (Glover & Mac Low 2007a, b) along with a simple treatment for CO chemistry (Nelson & Langer 1997; Glover & Clark 2012) with the AREPO code to simulate the chemical evolution of gas in the disc of the MW. Hydrogen is assumed to be present initially in fully atomic form, while carbon starts as C^+. The other species tracked in our chemical model (H^+, H_2, and CO) develop as the simulation progresses. The abundances by number relative to hydrogen for carbon and oxygen are chosen to be 1.4×10^{-4} and 3.2×10^{-4} (Sembach et al. 2000). The ratio of dust-to-gas is taken to be 0.01. The abundance of helium relative to hydrogen by number is 0.1. A spatially uniform ultraviolet radiation field is applied throughout the volume that has a strength of 1.7 times the Habing field (Habing 1968), similar to the MW's interstellar radiation field (Draine 1978). Gas self-shielding and dust shielding locally attenuate this field. We account for their effects using the TREECOL algorithm (Clark, Glover & Klessen 2012) with a maximum shielding length of 30 pc. Metal cooling of high-temperature gas assuming collisional ionization equilibrium is included with an assumed constant and spatially uniform gas metallicity of $1\,Z_\odot$ (Gnat & Ferland 2012;

---

[1] Any gas in our simulations that is predominantly molecular is generally too cold to excite the rotational energy levels of H_2, and hence is well-described by an adiabatic index $\gamma = 5/3$.

Walch et al. 2015). Ionization of atomic and molecular hydrogen from CRs is also included with a spatially uniform field that we set to have a strength of $3 \times 10^{-17}\,s^{-1}$ for H and twice that value for H_2 (Smith et al. 2014). In order to have consistent ionization fields between runs done with and without CRs we choose to not use a spatially varying CR energy density field to vary these ionization rates. By keeping ionization fields the same between all simulations, we can focus on the impact of CR pressure on the phase structure of the gas.

2.2.2 Thermal pressure support

The low temperature cooling provided by molecular gas and the gas self-gravity that we employ results in the collapse of gas to very high densities. Without a sink for dense gas or effective energetic feedback, a runaway collapse can occur. To model this collapse, which in the real ISM would result in the creation of dense molecular cores and stars that have sizes below the resolution of these simulations, we limit the pressure used in the Riemann solver. This approach has been used in AMR codes that place a limit on the finest mesh refinement level (e.g. Machacek, Bryan & Abel 2001). This approach was used in Simpson et al. (2016) where a minimum cell size was enforced. Here we have altered our approach to one better suited to a Lagrangian code. We begin with the assumption of a fixed mass resolution, and then compute the pressure of a Jeans stable cloud that can be resolved at this resolution.

We assume that a resolved cloud requires N_{cloud} gas cells. For a fixed cell mass m_t, this cloud will have a mass of $M_J = N_{\text{cloud}} m_t$. We estimate the mean density of a Jeans stable cloud of this mass with a mean temperature 10 times our chosen minimum allowed temperature (see below). This density is

$$\rho_{\text{max}} = \left(\frac{2\pi}{3 N_{\text{cloud}} m_t} \right)^2 \left(\frac{\pi \gamma (\gamma - 1) u_{\text{min}}}{G} \right)^3. \quad (3)$$

The pressure of gas at this density and temperature is

$$P_{\text{max}} = (\gamma - 1) u_{\text{min}} \rho_{\text{max}}, \quad (4)$$

where u_{min} is the specific internal energy of gas at 10 times the minimum temperature, i.e. 50 K. Once gas reaches the temperature floor, its pressure will scale linearly with density. However, as the gas density rises above ρ_{max}, the mass resolution is insufficient to resolve a Jeans stable cloud at that mean density. We therefore modify the fluxes computed in the Riemann solver to account for this issue.

We employ a pressure limit in the Riemann solver that follows a polytropic equation of state. For gas densities of ρ_i in cell i a minimum limit is imposed for both the left and right states in the Riemann solver:

$$P_{\text{floor,i}} = P_{\text{max}} \times \left(\frac{\rho_i}{\rho_{\text{max}}} \right)^\alpha. \quad (5)$$

The sound speed squared also has a floor of

$$c_{s,i}^2 = \gamma \frac{P_{\text{floor,i}}}{\rho_i}. \quad (6)$$

For $\alpha = 1$, the Riemann pressure floor would just be the cell's pressure, and the Riemann solver computes fluxes normally. To limit the flux of mass into dense clouds, we adopt $\alpha > 1$. The purpose of this study is not the detailed evolution of molecular clouds, we instead want to understand the overall phase structure of the ISM, including the hot and warm phases, where this numerical issue is not present, so modifying the pressure flows within dense clouds is an acceptable compromise. We have adopted a value of
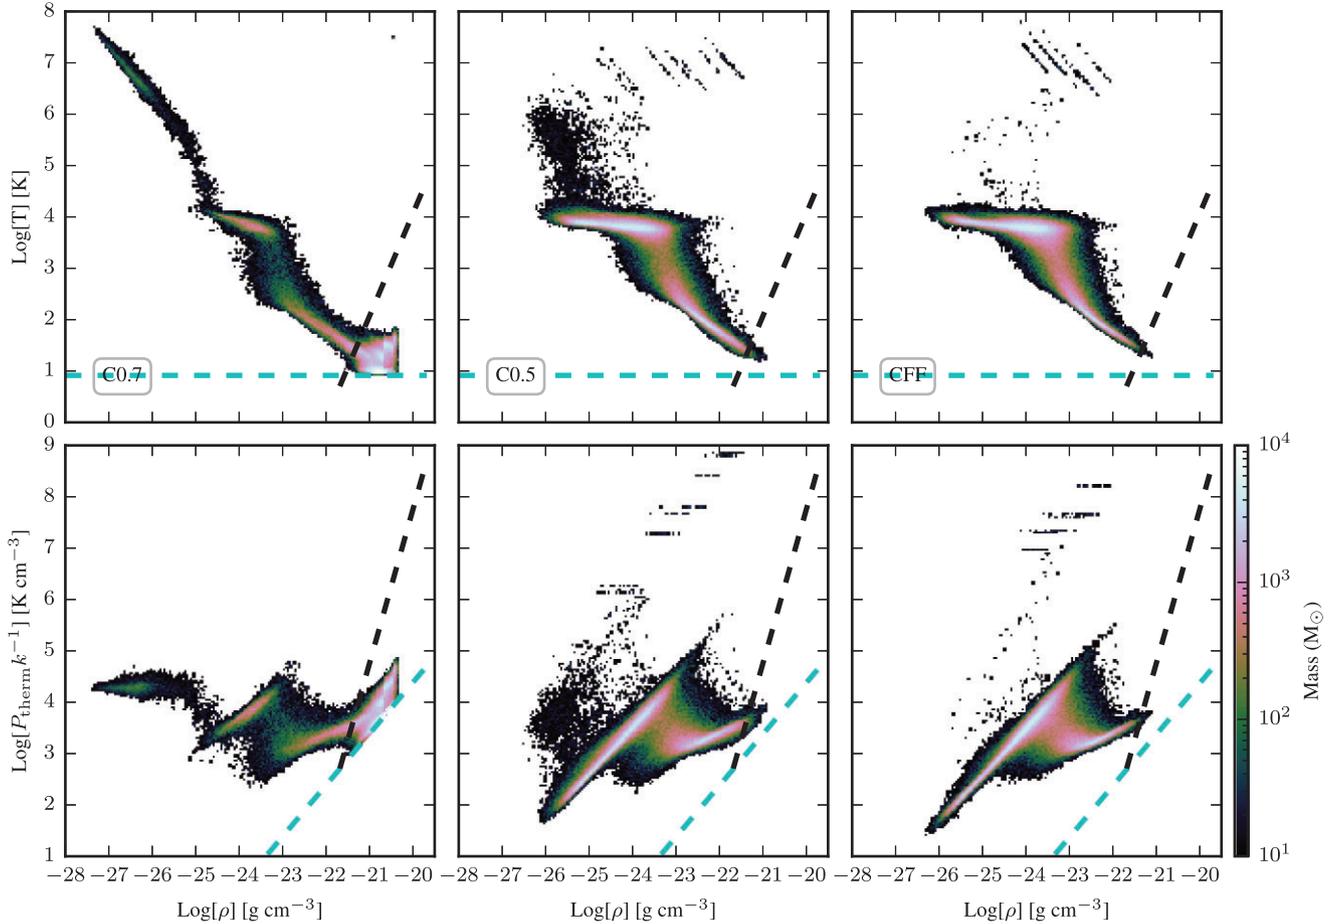
Figure 1. Density, temperature, and pressure of gas within 1 kpc of the mid-plane after 100 Myr of evolution for three representative runs with magnetic fields and CRs. Models with differing SN explosion placement are shown (see Table 1). The fraction of SN explosions placed in dense gas increases from left to right (see Section 3.1 for details). The intensity of the colour corresponds to the cumulative mass at each parameter-space coordinate: gas temperature versus mass density (top) and thermal pressure versus mass density (bottom). The density where the Riemann pressure floor activates is shown with a black dashed line and the pressure and temperature limits from the minimum allowed specific thermal energy is shown with cyan dashed lines.

$\alpha = 3$, which we found halted the collapse of gas to densities greatly exceeding ρ_{max} and did not affect the properties and evolution of gas at densities $\leq \rho_{max}$. Fig. 1 shows the pressures of gas using this model.

For our fiducial resolution, the target cell mass is $10 \, M_\odot$ and the minimum allowed temperature is 5 K. If N_{cloud} is 10^3, the Jeans mass M_J is $10^4 \, M_\odot$. The maximum stable density of a cloud of this mass at ten times the minimum allowed temperature is $\rho_{max} = 3.1 \, M_\odot \, pc^{-3}$, or $2.1 \times 10^{-22} \, g \, cm^{-3}$. In cells greater than this density, our model adjusts pressure gradients to push outward against the strong force of self-gravity that pulls gas inward. The sound speed criterion is necessary to ensure time-stepping remains Courant limited in these regions.

We have framed the motivation for this model in terms of the Jeans stability of a gravitating cloud by framing it in terms of the mass of the cloud ($N_{cloud} m_t$) and the specific energy of the cloud (u_{min}). However, the parameters N_{cloud} and u_{min} in equation (3) are degenerate with each other and the only parameter we have introduced in this model that has actual utility in the code is ρ_{max}. Ultimately, it is the value of ρ_{max} that is germane to the behaviour of the gas. We have performed tests adjusting ρ_{max} at fixed m_t in Appendix A. The evolution and properties of our models will be discussed in further sections, but Appendix A shows that the pressure floor model adjusts the maximum gas density reached in the simulation while preserving the overall phase structure of the medium.

2.2.3 Gravity

Gravity is included from both a static analytic potential representing a stellar disc and from the gas self-gravity.

The analytic stellar potential ϕ is a function of height h and is related to the initial gas density and gas fraction via Poisson's equation:

$$\nabla^2 \phi = 4\pi G \rho_0(h) \times \left(f_g^{-1} - 1 \right). \tag{7}$$

The function $\rho_0(h)$ is the initial gas density given by equation (1).

Gas self-gravity is computed with a tree-based algorithm with an adaptive softening length for gas cells that is related to a cell's volume by $\propto V^{1/3}$ (Springel 2010). We impose a minimum softening length of

$$\epsilon_{min} = \left(\frac{m_t}{\rho_{max}} \right)^{1/3} \tag{8}$$

For the fiducial model simulations the value of $\epsilon_{min} = 1.48$ pc. The tree algorithm utilizes an Ewald summation to capture long range gravitational forces at infinite distances across the periodic x and y boundaries (Springel et al. 2021).

2.2.4 Mesh refinement and resolution

Gas cells can be split or combined following a set of refinement and derefinement criteria that are dependent on the cells' masses,
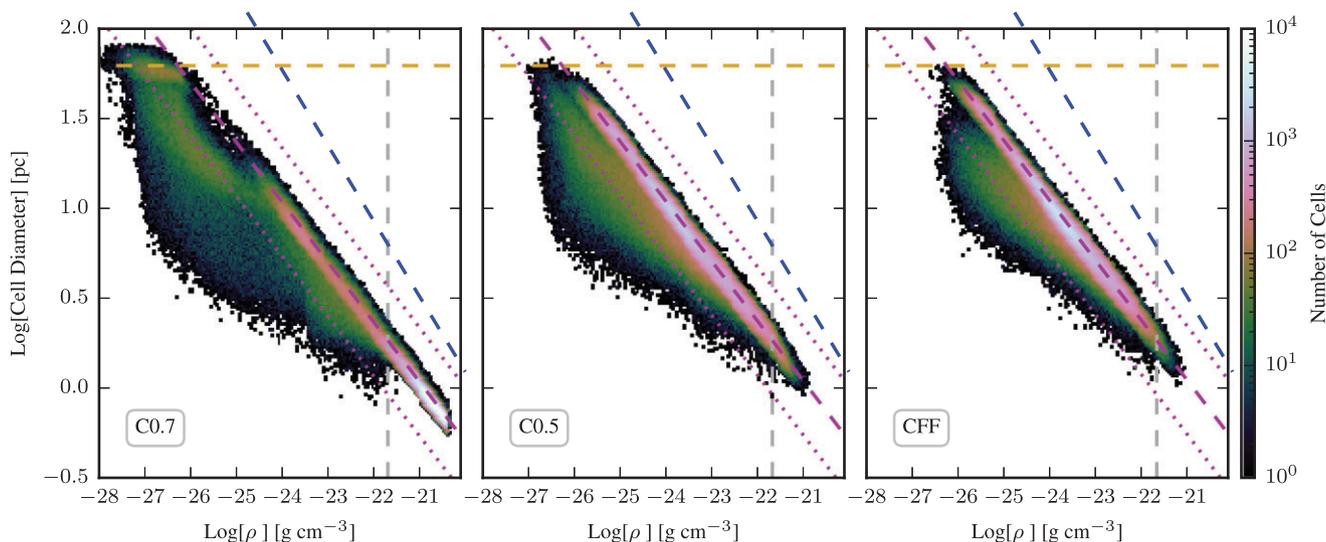
Figure 2. Resolution elements within 1 kpc of the mid-plane for three representative runs after 100 Myr of evolution that include magnetic fields and cosmic rays, and SN explosions. The 'cell diameter' (defined as $(6V_i/\pi)^{(1/3)}$ where V_i is the cell volume) is plotted versus the cell gas density. The colour intensity indicates the number of cells at each cell diameter and density. The magenta lines trace constant cell mass at the target gas mass $m_t = 10\,M_\odot$ (dashed) and at $2\times$ and $0.5 \times m_t$ (dotted). The diameter corresponding to the maximum cell volume (yellow dashed line) and the density where the pressure floor activates (grey dashed line) are also shown. Twice the cooling radius of an SNR with an explosion energy of 10^{51} erg is shown as a function of background density (e.g. Draine 2011, $\propto (\rho/m_p)^{-0.42}$, blue dashed line).

volumes, and Jeans lengths. Cells are refined to maintain masses within a factor of two of the target cell mass m_t. Cells are also refined in order to limit the volume ratio between adjacent cells to be no more than 10. In cases where the cell's diameter is more than four times the Jeans length (following Truelove et al. 1997) or where the maximum pressure difference among neighbouring cells is more than 10 times the cell's pressure, refinement is also applied. This latter criterion is only applied in gas less dense than ρ_{max}. No cell is allowed to have a volume greater than 2.5×10^5 pc^3 (which corresponds to an effective radius of 31 pc) and no limit is set for the minimum allowed cell volume.

The gas pressure floor and minimum softening halts the collapse of dense gas and truncates the distribution of cell sizes. We define an average cell diameter as $(6V_i/\pi)^{1/3}$. Fig. 2 shows the distribution of cell diameters for a few sample simulation outputs. The cell radius at which the pressure floor becomes relevant is below 2 pc, so only a small fraction of the simulation volume is covered by these cells. However, these cells can contain a large amount of mass if numerous enough. Other scales of the ISM are well resolved. The scale height of the disc and the SN remnant (SNR) cooling radius at the volumetric average gas density are resolved by many cells unaffected by the pressure floor. The cooling radius at the mass-weighted average density is also a factor of ∼3 greater than the median cell radius.

Fig. 1 shows the phase space distribution that results from this Riemann pressure floor model. The dense gas supply is heavily influenced by the SN placement model, discussed in future sections, but the behaviour of gas with pressure below the Riemann pressure floor (shown by the cyan line in Fig. 1) is limited by this model. This gas consists of dense clouds as can be seen by eye in Fig. 4. We do not address the internal structure or evolution of these clouds as this polytropic model is not intended to capture the physics that regulates those properties. We note that the deposition of feedback energy in these clouds will raise the pressure locally and thus make the pressure floor non-relevant for those affected cells. Appendix A describes in more detail how this model behaves with different model assumptions.

2.3 Cosmic ray energy and transport

CRs are modelled as a second fluid with an adiabatic index of $\gamma_{CR} = 4/3$ as described by Pfrommer et al. (2017a). This fluid tracks and evolves the energy density of CRs and couples the CR fluid to the thermal gas fluid via the CR pressure. We additionally include CR transport in the form of anisotropic diffusion as described by Pakmor et al. (2016b). This model assumes a spatially and temporally constant diffusion coefficient κ along the direction of the magnetic field in the cell and zero in all other directions. We use a fiducial value of κ of 10^{28} cm^2 s^{-1} (and vary it from 10^{27} to 10^{29} cm^2 s^{-1} to address the uncertainties associated with this choice). The CR energy density in our simulations evolves according to this equation:

$$\frac{\partial \varepsilon_{CR}}{\partial t} + \nabla \cdot [\varepsilon_{CR} \boldsymbol{v} - \kappa \boldsymbol{b}(\boldsymbol{b} \cdot \nabla \varepsilon_{CR})] = -P_{CR} \nabla \cdot \boldsymbol{v} - \Lambda_{CR} + \Gamma_{CR}, \quad (9)$$

where ε_{CR} is the energy density of CRs, $\boldsymbol{v}$ is the gas velocity, $\boldsymbol{b}$ is a unit vector in the direction of the local magnetic field, P_{CR} is the CR pressure, Γ_{CR} is CR energy injected in SNe, and Λ_{CR} is the loss of CR energy from hadronic and Coulomb processes.

These Coulomb and hadronic processes are dependent on the energy spectrum of CRs. Unlike schemes that dynamically follow the CR energy spectrum on top of the MHD (Girichidis et al. 2020, 2022), the two-fluid model we use does not self-consistently track the spectrum of CRs, and therefore, we employ the assumption of an equilibrium spectrum that balances CR injection and loss (Enßlin et al. 2007; Pfrommer et al. 2017a). We assume that all Coulomb losses and one-sixth of hadronic losses are thermalized, resulting in a transfer of CR energy to thermal energy. The remaining hadronic losses (5/6) are assumed to be directly lost to gamma-rays and neutrinos, which are radiated away for the optically thin conditions of the ISM.

For the adopted typical spectral index of shock-accelerated CRs of $\alpha_{inj} = 2.2$ (Caprioli, Haggerty & Blasi 2020; Haggerty & Caprioli
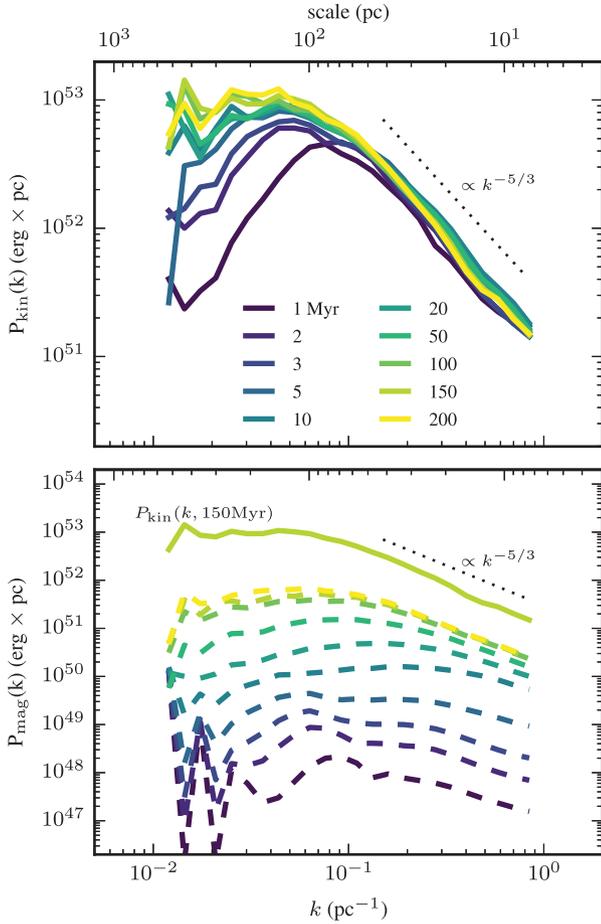
Figure 3. Top: Power spectra of kinetic energy within 250 pc of the box mid-plane at different times ranging from 1 to 200 Myr in simulation C0.5, which contains CRs and magnetic fields. The Kolmogorov turbulence slope of $-5/3$ is indicated by a dotted line. Bottom: Power spectra of magnetic energy in the same volume shown with dashed lines. For comparison, the power spectrum of kinetic energy at 150 Myr from the top panel is shown. The power-law slope of $-5/3$ is also indicated.

2020), the Coulomb loss rate is taken from (Pfrommer et al. 2017a):

$$\Lambda_{\mathrm{Coul}} = 2.78 \times 10^{-16} \left(\frac{n_e}{\mathrm{cm}^{-3}} \right) \left(\frac{\varepsilon_{\mathrm{CR}}}{\mathrm{erg\ cm}^{-3}} \right) \mathrm{erg\ s}^{-1} \mathrm{cm}^{-3} \quad (10)$$

and the hadronic loss rate by

$$\Lambda_{\mathrm{hadr}} = 7.44 \times 10^{-16} \left(\frac{n_e}{\mathrm{cm}^{-3}} \right) \left(\frac{\varepsilon_{\mathrm{CR}}}{\mathrm{erg\ cm}^{-3}} \right) \mathrm{erg\ s}^{-1} \mathrm{cm}^{-3}. \quad (11)$$

These rates are used during the running of the simulation to provide sources and sinks to the CR and thermal fluids as we have described.

2.4 SN injection models

SN explosions are modelled as discrete events depositing 10^{51} erg of energy in the simulation when they occur. They are deposited in the gas by injecting energy into the 32 closest cells surrounding the site selected for the explosion. This site can either be the centre of a gas cell with a desired property or a point in space selected by other criteria.

The SN energy that is injected is split between thermal energy and CR energy. We choose to inject 10 per cent of the SN energy in CRs and 90 per cent in thermal energy. This value is consistent with

previous models (e.g. Simpson et al. 2016) and some observations (e.g. Morlino & Caprioli 2012) but may be on the high side of acceptable values if magnetic obliquity-dependent acceleration efficiencies are integrated around the entire supernova remnant (Caprioli & Spitkovsky 2014; Pais et al. 2018). Both types of energy are injected proportionally to a cell's volume in the SN region. In models without CRs, all 10^{51} erg of energy is deposited in thermal energy. The chemical composition of cells are not adjusted during the injection.

We use two methods for SN explosion placement. The first method (called FF) was used in Simpson et al. (2016) and selects SN explosion locations based on a local SFR computed from the local dynamical free-fall time:

$$\mathrm{sfr}_i^{\mathrm{ff}} = \epsilon \frac{m_i}{t_{\mathrm{ff},i}}, \quad (12)$$

where m_i is the mass of cell i, ϵ is a star formation efficiency (we adopt 1 per cent), and $t_{\mathrm{ff},i}$ is the dynamical free-fall time within the gas cell:

$$t_{\mathrm{ff},i} = \sqrt{\frac{3\pi}{32 G \rho_{\mathrm{b},i}}}. \quad (13)$$

The dynamical free-fall time is a function of the baryon density with the cell $\rho_{\mathrm{b},i}$, which is the sum of the gas density ρ_i and the stellar mass density given by the background potential profile.

We assume the rate of SN explosions in each cell is related to $\mathrm{sfr}_i^{\mathrm{ff}}$ by the factor of 1.8 explosions per 100 $\mathrm{M}_\odot$ of newly formed stars. The probability of a supernova explosion at the position of cell i in time step Δt is therefore

$$p_i^{\mathrm{ff}} = \mathrm{sfr}_i^{\mathrm{ff}} \times \frac{1.8\ \mathrm{SNe}}{100\ \mathrm{M}_\odot} \times \Delta t, \quad (14)$$

where we have assumed that $\Delta t \ll t_{\mathrm{ff},i}$. Indeed, this is the case as seen in Fig. A3 in Appendix B where we see that t_{ff} is well above the maximum time-step of 0.1 Myr. In dense gas where t_{ff} is shorter, the time-steps are typically much less than 0.1 Myr in order that the evolution of the thermal gas is Courant limited.

The second method of SN placement (called MIX) adds a random-placement component to the FF model. This component reflects SN explosions that do not occur in their natal molecular cloud, such as those arising from run-away massive stars, or that occur in star-forming environments where dense gas has already been cleared away by prior SN explosions, photoionization or winds from Wolf–Rayet stars. For this method, we select a fraction of randomly placed SN explosions, f_{rand}, and divide the total number of SN explosions between the two modes. The total number of SN explosions is fixed according to the empirical Kennicutt–Schmidt relation (Kennicutt 1998) that gives the surface density of star formation $\dot{\Sigma}_{\mathrm{KS}} = 2.5 \times 10^{-4} \Sigma_{\mathrm{g}}^{1.4}$ ($\mathrm{M}_\odot\ \mathrm{kpc}^{-2}\ \mathrm{yr}^{-1}$), as a function of gas surface density Σ_{g} ($\mathrm{M}_\odot\ \mathrm{pc}^{-2}$). The number of randomly placed SN explosions in time step Δt is

$$N_{\mathrm{rand}} = \dot{\Sigma}_{\mathrm{KS}} \times L_{\mathrm{box}}^2 \times \frac{1.8\ \mathrm{SNe}}{100\ \mathrm{M}_\odot} \times f_{\mathrm{rand}} \Delta t. \quad (15)$$

SN explosions injected in this mode are placed with uniform random-ness in the horizontal $x - y$ plane. Their vertical z position is drawn randomly from a distribution function that follows equation (1), so that events are concentrated in the mid-plane. These randomly placed SN explosions therefore represent a population of core-collapse events from massive stars, however, explosions from Type Ia SN would also have no spatial correlation with dense gas, but would also not depend on the Kennicutt–Schmidt relation.

The remaining SN explosions in MIX are placed according to local gas properties in a manner similar to FF. Each cell is given a local SFR for this mode:

$$\mathrm{sfr}_i^{\mathrm{mix}} = \mathrm{sfr}_i^{\mathrm{ff}} \times \frac{\dot{\Sigma}_{\mathrm{KS}} L_{\mathrm{box}}^2}{\sum \mathrm{sfr}_i^{\mathrm{ff}}} \times (1 - f_{\mathrm{rand}}) \qquad (16)$$

and the probability p_i^{mix} of an SN explosion occurring at the position of cell i is computed in the same way as equation (14).

By normalizing the local SFR in this way, the MIX model fixes the global SFR rate to the Kennicutt–Schmidt rate, while maintaining the relative spatial differences in the SFR used to place SNR in the non-random mode due to variations in gas density. In the MIX model, the global SFR is adjusted over time as gas flows out of the simulation box at the outflow boundaries and Σ_{g} decreases. In the FF model, the global SFR varies due to changes in local gas properties, the effect of which is reflected in the global rate because the global rate is the sum of all the local SFRs.

The net result of increasing f_{rand} is that more SN energy will be coupled to low density gas on injection. Given our Lagrangian mesh refinement scheme, this means the injection zones in models with larger f_{rand}, while always comprising 32 cells and containing approximately 320 M$_\odot$ of gas, will occupy a larger volume. However, the mesh resolution requirement for capturing the cooling of SN energy deposited on a mesh decreases with decreasing background density (e.g. Simpson et al. 2015). Also, the cooling time in low density gas in runs where it dominates the volume is typically longer than the simulation time (see Fig. A3).

For consistency with the assumed SFR we compute locally, we also employ a step of draining star forming gas from cells and re-injecting it onto the mesh to account for gas loss from star formation. This is done instantaneously within a time step, but a very small amount of gas is cycled in this way, reflecting the low efficiency of star formation on the scales we consider here. At each time step and in each cell i, we estimate the amount of gas mass that goes into stars as $\mathrm{sfr}_i \times \Delta t$. For the FF model, sfr_i is simply $\mathrm{sfr}_i^{\mathrm{ff}}$ and for the MIX model, it is $\mathrm{sfr}_i^{\mathrm{mix}}/(1 - f_{\mathrm{rand}})$. If this mass is less than the mass of the cell, it is removed. Once all star forming mass is drained, the total amount of gas is redistributed on the mesh. In this redistribution phase, each cell receives an amount of gas proportional to $\rho_0(h_i)(1/f_{\mathrm{g}} - 1)V_i$, which is the mass of stars implied by our assumed potential, gas fraction, and the cell volume. When mass is drained the specific CR energy is kept constant, resulting in a loss of CR energy. When gas is injected, the specific CR energy is changed to keep the CR energy constant. Thus, this cycling of star-forming gas is a net sink for CRs but not a source of CR heating. The magnetic energy is kept fixed during both steps.

Our SN injection model does not directly inject kinetic or magnetic energy, but the pressure forces caused by the injection of thermal energy induces gas motions that results in a turbulent medium, which in turn, amplifies the seed field. Fig. 3 shows the development of the turbulent spectrum of kinetic energy that has a Kolmogorov-like slope from scales of pcs to tens of pcs and saturates very early in the simulation. The spectrum of magnetic energy (also shown) grows more slowly but also reaches saturation by 100 Myr.

3 HOW CRS MEDIATE THE EFFECT OF SN EXPLOSIONS ON ISM STRUCTURE

In this section we focus on how the structure of the ISM changes when different models for SN placement are applied. The models for SN placement that we test vary the number of SN explosions placed

Table 1. List of simulations.

Simulation	f_{rand}	MHD	CR	κ (cm^2 s^{-1})
HFF	–	no	no	–
CFF	–	yes	yes	10^{28}
H0.1–H1.0*	0.1–1.0	no	no	–
C0.1–C1.0*	0.1–1.0	yes	yes	10^{28}
CFF-κ27-29**	–	yes	yes	10^{27}–10^{29}
C0.5-κ27-29**	0.5	yes	yes	10^{27}–10^{29}
CFF-adv†	–	yes	yes	0
C0.5-adv†	0.5	yes	yes	0

Notes. *These simulations vary f_{rand} from a value of 0.1 to 1.0 as indicated in the simulation label.
**These simulations vary the CR diffusivity κ from a value of 10^{27} cm^{2}s^{-1} to 10^{29} cm^{2}s^{-1}. The value of Log(κ) is indicated in the simulation label.
†These simulations include the CR fluid, but only allow it to advect with the thermal gas; the diffusion term from equation (9) is not included (i.e. $\kappa = 0$).

in high-density gas and the number of remnants placed randomly in the mid-plane. We explore how varying the fraction of SN energy coupled to low density gas (in models with high numbers of randomly placed SN explosions) increases the volume filling factor of hot gas ($\geq 2 \times 10^4$ K) and we will then explore how the presence of CR pressure modifies this effect. Table 1 summarizes the simulations run to model these effects.

3.1 Varying the placement of SN explosions

Fig. 4 shows a snapshot of the state of the gas after 100 Myr of evolution in a representative sample of our models with different SN explosion placement. It shows the central kpc of each simulation volume that includes the multiphase mid-plane gas. Fig. 4 shows that gas in the mid-plane becomes more porous as models include higher fractions of randomly placed SN explosions. In models with predominantly random SN explosions, a 'thermal runaway' occurs in the mid-plane, where the coupling between SN energy and dense gas is inefficient and where gas self-gravity and radiative cooling cause clouds to become denser and denser. Subsequent explosions tend to occur in hot, rarefied gas as the hot phase occupies most of the volume (e.g. McKee & Ostriker 1977). This heats gas further, and the end result is that the volume fraction filled by gas with temperatures in excess of 2×10^5 K exceeds 90 per cent. In models with high f_{rand}, the mid-plane mass becomes dominated by a population of dense, self-gravitating clouds. The size of these clouds is regulated by the Riemann pressure floor that acts to keep them Jeans stable, i.e. prevents them from collapsing further.

In models with low f_{rand}, SN energy couples to dense gas more efficiently because of the placement of more SN explosions in dense gas. This 'peak driving' prevents gas from collapsing to dense clouds, consistent with previous studies (e.g. Walch et al. 2015). In these models, 'warm gas' (with temperatures between 300 K and 2×10^5 K) dominates the volume.

Fig. 5 shows how this process proceeds over time through the evolution of mid-plane properties. In the MIX models, the global SFR is fixed to the Kennicutt–Schmidt rate of the box. All the MIX models begin with the same SFR, but as time goes on, and as gas is lost from the box through outflows (the topic of a companion paper), the SFR drops because the gas surface density drops.

In the FF model, the SFR is the sum of all the local rates given by equation 12), i.e. determined only by the local gas density and
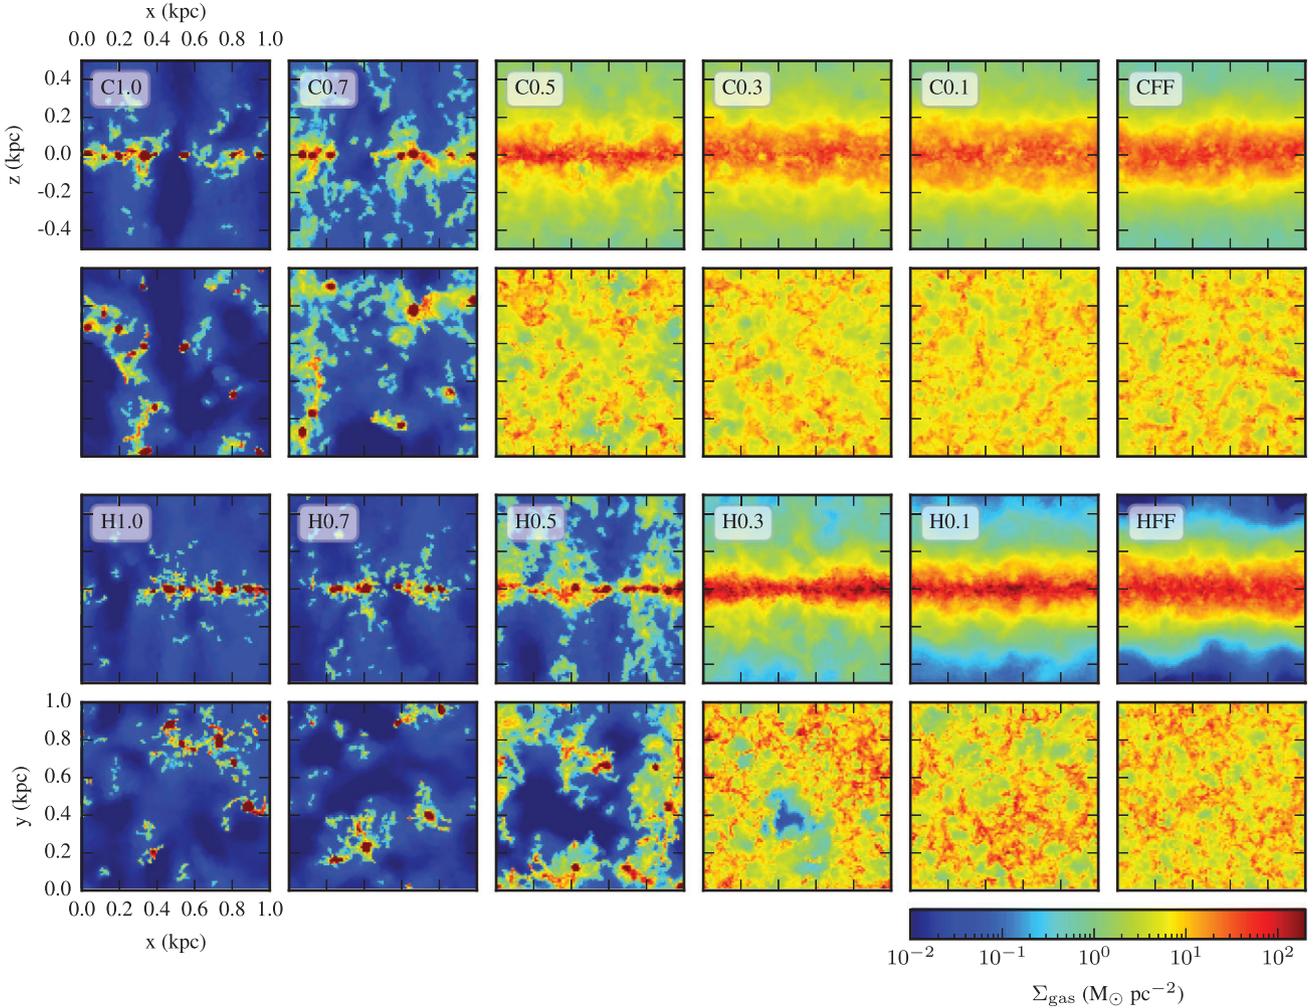
Figure 4. Projections of gas number density of the central square kiloparsec of models with CRs (top two rows) and without CRs (bottom two rows). In each pair of rows, two orthogonal projections are shown: edge-on (top) and face-on (bottom). From left to right, simulations with varying SNR injection models are shown, and are in order MIX models with $f_{\rm rand}$ = 1, 0.7, 0.5, 0.3, and 0.1, and then the FF model.

free-fall time, instead of being fixed to a global rate. This gives an initial rate close to the Kennicut–Schmidt rate, but due to the initial radiative cooling of the mid-plane gas, a modest initial starburst occurs that gives a 50 per cent enhancement in the SFR. By 50 Myr, this burst is arrested and an equilibrium state is reached in the FF model that balances radiative cooling, SN energy injection, and mass-loss through outflows, resulting in a near constant SFR and median mid-plane gas properties. The volume-weighted median mid-plane gas density and temperature (defined as the density and temperature that bisect the volume-weighted distributions of these quantities), are essentially constant ($0.3\,{\rm cm}^{-3}$ and 7000 K), and the height of the central gas layer (defined as the height containing half the original gas mass) is also essentially constant in FF.

In the MIX models, the mid-plane evolution proceeds differently and depends strongly on $f_{\rm rand}$. There appear to be two attractor solutions to which MIX models will evolve: one is the thermal runaway solution, which occurs when $f_{\rm rand} \geq 0.6$ and the other is the peak driving solution, which occurs when $f_{\rm rand} \leq 0.5$. In the thermal runaway state, mid-plane gas evolves to gas densities below $10^{-3}\,{\rm cm}^{-3}$ and temperatures in excess of 10^7 K.

The vertical distribution of gas also bifurcates into two states. With large $f_{\rm rand}$ the mass in the mid-plane collapses to pressure-floor supported clumps as seen in Fig. 4 that then results in a gas layer

with a height given by the ratio of the Jeans mass we resolve ($M_{\rm J}$) and the density ($\rho_{\rm max}$) where the pressure floor kicks in:

$$h_{\rm min} = \left(\frac{3}{4\pi} \frac{M_{\rm J}}{\rho_{\rm max}} \right)^{1/3} \sim 12\ {\rm pc}. \qquad (17)$$

With a small $f_{\rm rand}$, the height of the gas layer extends above 100 pc. In the MIX models with a small $f_{\rm rand}$, the height of the gas layer does not reach an equilibrium, unlike FF, because the global SFR is larger and does not respond to instantaneous changes in the mid-plane; it adjusts downward only when gas is lost through the outflow boundaries. This results in a constant puffing up of the mid-plane gas layer.

The closest model to an 'intermediate' state between the two attractor solutions is the $f_{\rm rand} = 0.6$ model (pink lines in Fig. 5). It has a slower rise in its gas scale height in the first 100 Myr and a mid-plane gas density close to the peak driving runs at 50 Myr. However, this state is not stable, and after 100 Myr, it evolves to the thermal runaway case.

In later sections, we will discuss how the evolution shown in Fig. 5 varies when CR energy transport is altered, or when CRs are excluded all together. In most of our models, the two general solutions described here appear to be the stable states that the simulated ISM is driven to; what varies when the energy partition or the CR transport is
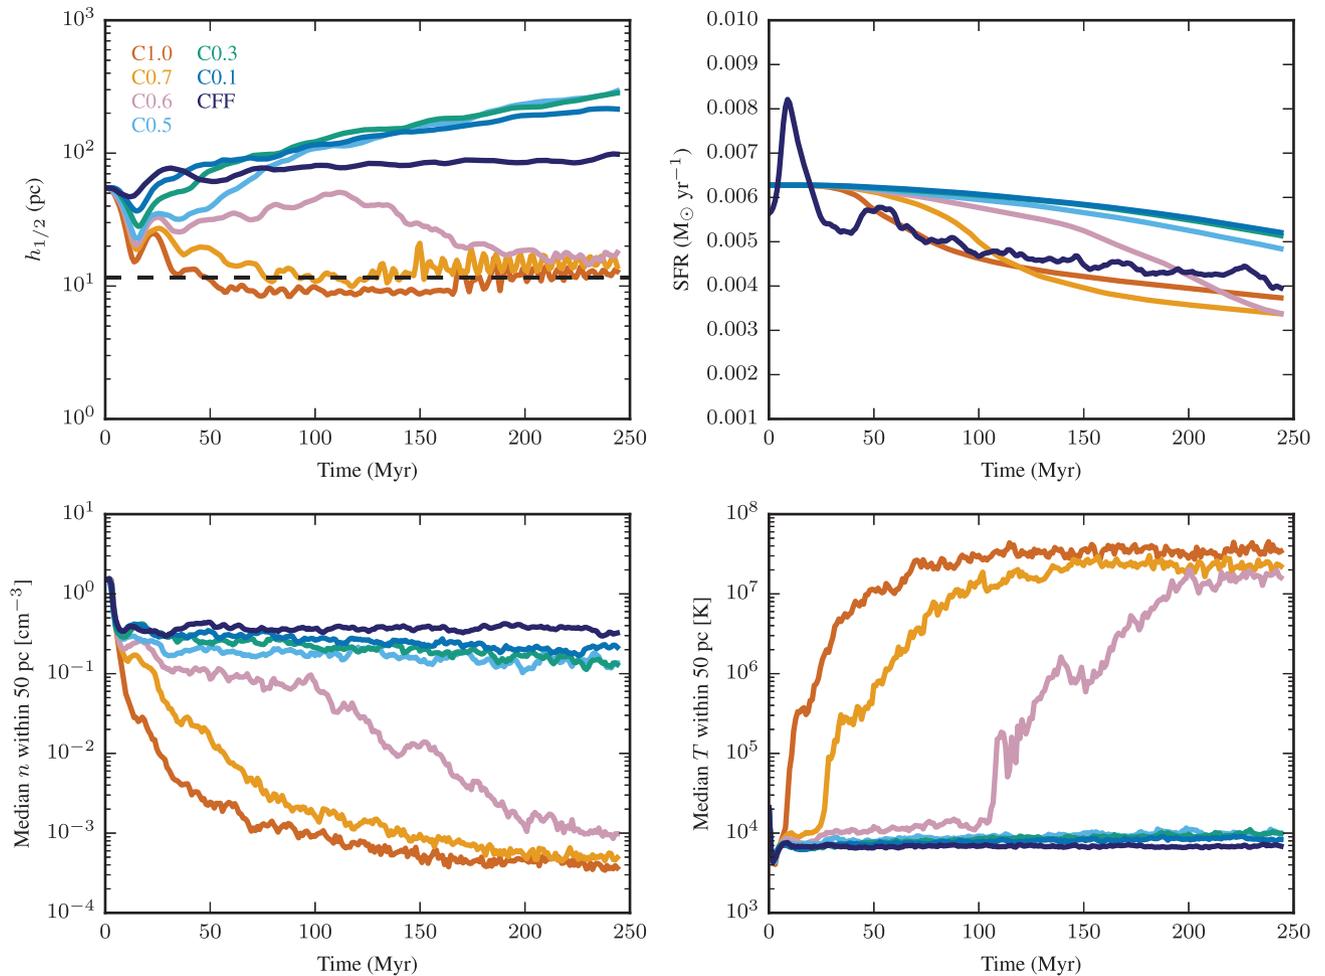
Figure 5. Evolution of mid-plane properties for simulations with CRs. Top left-hand panel: Height containing half the original gas mass within the box. The minimum radius of a Jeans stable cloud assumed by the pressure floor model is shown by a dashed line. Top right-hand panel: The total star formation rate of the box, which seeds the SNR injection rate. Bottom left-hand panel: The volume-weighted median gas number density within 50 pc of the mid-plane. We compute this median value with respect to cell volume, i.e. this is the density at which half the volume is at lower densities and half at higher. Bottom right-hand panel: The volume-weighted median gas temperature within 50 pc of the mid-plane. We adopt the same volume-weighted median definition as in the number density panel.

altered is the timing of the evolution and the tipping-point of energy injection that divides one state from another.

Fig. 6 shows the differences in the distributions of mid-plane gas between the peak-driving and thermal run-away states. The distribution in Log(n) (n is the gas number density) follows a skewed log-normal distribution in the peak driving case, with the skewness extending to the low density side. The distribution in the thermal runaway case follows approximately a power law in Log(n), with three bumps: one at densities above ρ_{max} due to pressure-floor supported gas clumps; one at Log(n) =0, where the peak driving distributions peak; and one at Log(n) =−3 where the distribution turns over. We note that the volume-weighted average of n and the volume-weighted average of Log(n) occur at different n, with the average in Log(n) occurring at lower n.

Fig. 6 also shows the vertical distribution of gas density. Simulations in the thermal runaway case have a rapid fall off in density beyond 11 pc, approximately the radius of a Jeans stable cloud at the Riemann pressure floor. For simulations in the peak driving case, the gas density falls off much more gradually. The typical e-folding height of gas (the height where the gas surface density drops by 1/e), is less than 20 pc in the first case and between 50 and 150 pc in the second case. We note that our simulations do not include a

circumgalactic medium (CGM) component that the disc medium can interact with, which would likely impact these vertical profiles at larger heights.

## 3.2 Mid-plane energy partition

The phase structure differences we have described correlate with different energetic balances within the ISM. Fig. 7 shows the spatial distribution of CR energy and temperature, along with the ratio of thermal energy to the three other energies in the simulation, magnetic, CR, and kinetic. We define these ratios as β_{mag}, β_{CR}, and β_{kin}:

$$\beta_{kin} = \frac{\varepsilon_{therm}}{\frac{1}{2}\rho v^2}, \tag{18}$$

$$\beta_{CR} = \frac{\varepsilon_{therm}}{\varepsilon_{CR}}, \tag{19}$$

$$\beta_{mag} = \frac{\varepsilon_{therm}}{B^2/(8\pi)}, \tag{20}$$

where ε_{therm} is the thermal energy density. A β value of 1 indicates a cell with equal amounts of thermal energy and the non-thermal energy under consideration.
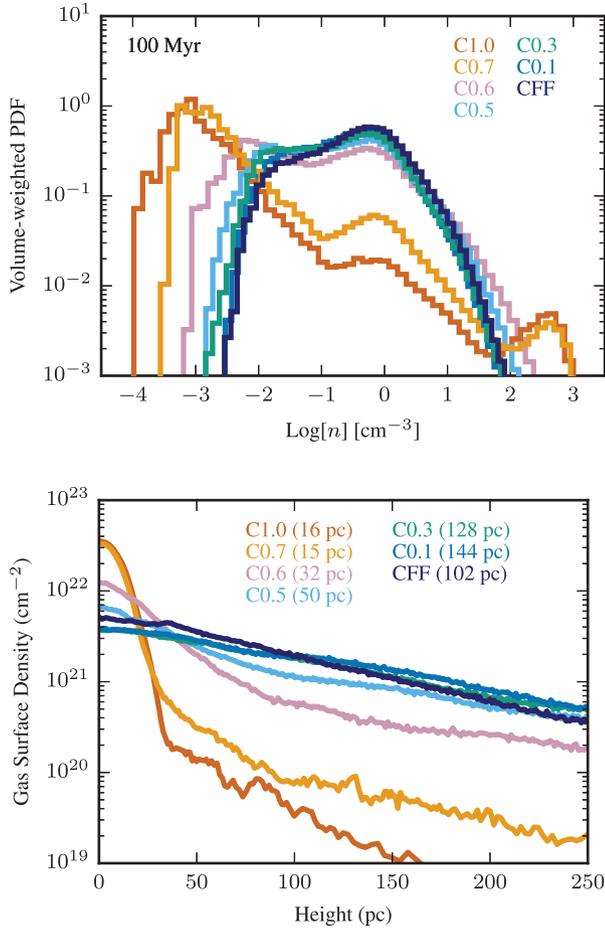
Figure 6. Top: Volume-weighted probability density functions (PDFs) of gas density within 50 pc of the mid-plane for simulations with cosmic ray diffusion after 100 Myr of evolution. The PDFs shown are of the logarithms of the gas number density, Log(n). Bottom: Surface density of gas in the vertical direction for models with cosmic ray diffusion after 100 Myr of evolution. The simulation volume above and below the mid-plane are combined and the mid-plane is at height 0. The height where the gas surface density drops by $1/e$ from the peak value at the mid-plane is indicated for each run.

In all models, the CR energy density is quite smooth throughout the mid-plane, with the exception of the location of very young SN explosions. The thermal gas energy varies more, so the spatial variation of β_{CR} is set primarily by the thermal energy. In the peak driving runs, most of the volume is dominated by CR energy and only in regions with young SNRs does thermal energy dominate due to our prescribed injection ratio of CRs. In thermal runaway runs, most of the volume is dominated by thermal energy and CRs only dominate in cold, dense clouds.

The overall energy density of CRs does vary between high and low f_{rand} runs, with a higher CR energy density level found in the mid-plane in the peak-driving runs with lower f_{rand}. With most of the volume of the peak-driving runs being filled with warm ($\sim 10^4$ K) gas rather than hot ($\sim 10^7$ K) gas (as in the thermal-runaway case), on large scales the CR energy dominates the volume in the peak-driving case, unlike the thermal-runaway case.

High magnetic energy density is largely confined, in both types of model, to only the densest gas and is therefore subdominant in most of the mid-plane volume. In models with large peak driving, the ratio of thermal to turbulent energy can vary on quite short length scales. In models with large random driving, the thermal energy is

so great and the gas density so low that the thermal energy exceeds the turbulent energy.

In the thermal-runaway case (e.g. C0.7), the thermal energy dominates on large scales and it is only in the dense clouds where magnetic and kinetic energies can dominate. These clouds are strongly self-gravitating, but are subject to the pressure limiter, so their internal properties should be regarded in that light. In the peak-driving case (e.g. C0.5), the kinetic and thermal energies are nearly in equipartition on large scales, but have variations on smaller scales following variations in the local thermal temperature and density which span a smaller dynamic range.

In all runs, the magnetic β_{mag} is large (>10) in most of the volume. Pockets where it is less than unity correlate with areas of dense gas, however in peak-driving runs (e.g. C0.5), there is a larger fraction of the volume with $\beta_{mag} \sim 1$, found in ~ 1000 K gas.

The trends on large scales are summarized in Fig. 8. Models of the ISM at the solar circle have long indicated that the different pressures in the ISM, thermal, kinetic, CR, and magnetic, are nearly in equipartition at the mid-plane (Boulares & Cox 1990). The thermal-runaway models have larger densities of thermal energy in the mid-plane; however, the other energetic components of the ISM are not equivalently elevated, so that no other energy component is more than 10 per cent of the thermal energy density. This contrasts with the peak-driving runs in which the kinetic energy density is within a factor of a few of the thermal energy density and the CR energy density is typically a factor of 2–3 higher than the thermal energy density. The magnetic energy density is below 10 per cent of the thermal energy consistent with the saturation strength expected for a small-scale turbulent dynamo. By construction we lack cosmological (Marinacci & Vogelsberger 2016; Pakmor et al. 2017; Pfrommer et al. 2021) and disc amplification mechanisms (Shukurov et al. 2006) that could amplify the magnetic field beyond this strength.

3.3 Evolution of the CR energy

The β_{CR} parameter is determined by the evolution of the thermal gas, but also by the evolution of the CR fluid, which Fig. 7 shows does not have the same dependence on the phase structure of the gas as the thermal energy does. There are several source and sink terms for the CR energy (see equation 9) in our simulations. There is the injection of CR energy through our SN feedback model (Γ_{CR}); adiabatic changes of the CR fluid due to gas motions ($-P_{CR}\nabla\cdot v$); CR cooling from both hadronic and Coulomb losses (Λ_{CR}); and losses due to star formation that arise when we remove star-forming mass from a cell and keep the specific CR energy constant. In addition, CR energy can be lost through the outflow boundary, either advectively or diffusively.

Fig. 9 shows the evolution of the total CR energy in two representative simulations and the evolution of all the source and sink terms. These simulations create an outflow from the mid-plane. This is associated with a net expansion in the gravitationally stratified atmosphere and hence, a net adiabatic loss of CR energy. It is this work that is the dominant loss. In the model with the thermal runaway, the CR injection and adiabatic losses appear to be nearly in equilibrium, so CRs expand from the mid-plane almost as soon as they are injected. Due to the injection of CR energy in a low density medium, it almost immediately expands adiabatically. In models with peak driving, CR energy does not expand as readily and remains confined to the warm layer of gas for a longer period of time. This also results in a different amount of CR energy being lost across the outflow boundaries, with it being higher in the thermal runaway models.
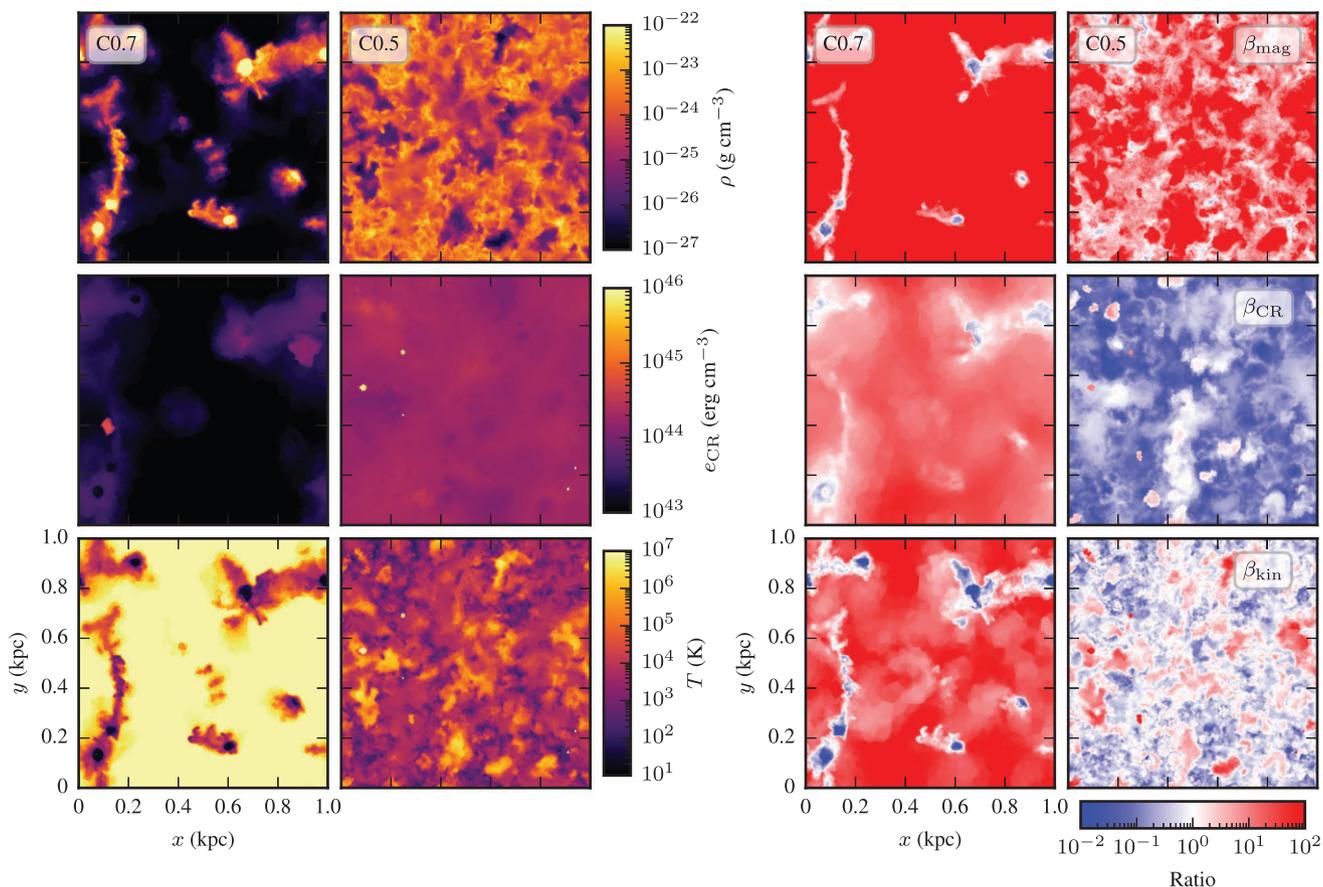
Figure 7. Thin projections of mid-plane gas properties in the C0.7 and C0.5 models after 100 Myr of evolution. The central 100 pc of the box are projected along the vertical axis showing a face-on view and the full width of the box (1 kpc × 1 kpc) is shown. Left-hand panels: the panels show, from top to bottom, the gas density; the CR energy density weighted by the gas density; and the gas temperature weighted by the gas density. Right-hand panels: these panels show the ratio of energy column density maps for different energetic components. From top to bottom, we show thermal energy density to magnetic energy density (β_{mag}); the thermal energy density to CR energy density (β_{CR}); and the thermal energy density to kinetic energy density (β_{kin}). White indicates regions of equal energy density and red thermally dominated regions.
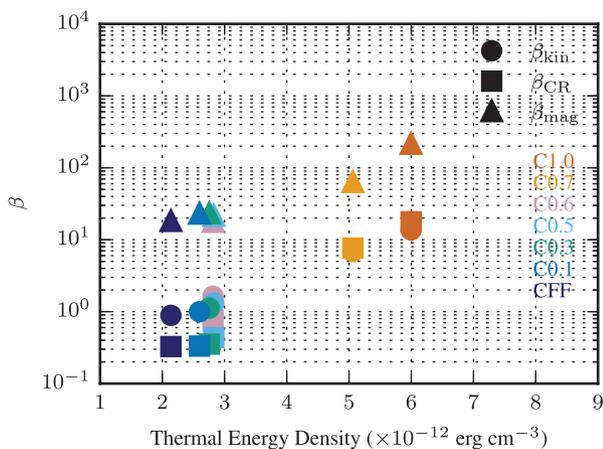

Figure 8. Energy ratios β_{kin}, β_{CR}, and β_{mag} of gas within 50 pc of the mid-plane. Quantities are averaged from snapshots between 95 and 104 Myr. The thermal energy density is defined as the sum of all cells' thermal energy divided by the total volume of the mid-plane region (1 kpc × 1 kpc × 100 pc). The ratios shown are the thermal-to-kinetic energy (β_{kin}, circles), the thermal-to-CR energy (squares), and the thermal-to-magnetic energy (triangles).

There are small differences in the amount of energy injected because there is a small difference in the SFR of these models, which sets this term. The losses due to gas removal by star formation are cumulatively small, amounting to at most a few SN explosions worth of energy and subdominant in all cases.

The CR cooling is similar in both models. However, it is less than the adiabatic losses in both cases. Thus, the CR energy contributes to the powering of outflows from the mid-plane. Also in both cases, the hadronic cooling losses dominate over the Coulomb losses by a factor of more than 2.5. While the thermal runaway model has more gas mass at high densities, which has a higher rate of CR-nucleon collisions, it occupies such a small volume that it only removes a small amount of CR energy. In the peak driving run, the CR cooling approaches the adiabatic losses and arises from the warm neutral phase of gas that occupies a large volume in this run.

In summary, the reason the mid-plane thermal energy density dominates over the CR energy density in the thermal runaway runs is the fragmentation of the mid-plane gas. This results in the dual effects of thermally heating the mid-plane gas to high temperatures and allowing for the more efficient outflow of CR energy deposited in the mid-plane, producing an overall drop in CR energy in the calculation. In the peak-driving case, the stability of the warm gas layer results in overall lower mid-plane thermal temperatures and the less efficient escape of CR energy through $P\mathrm{d}V$ work. The only
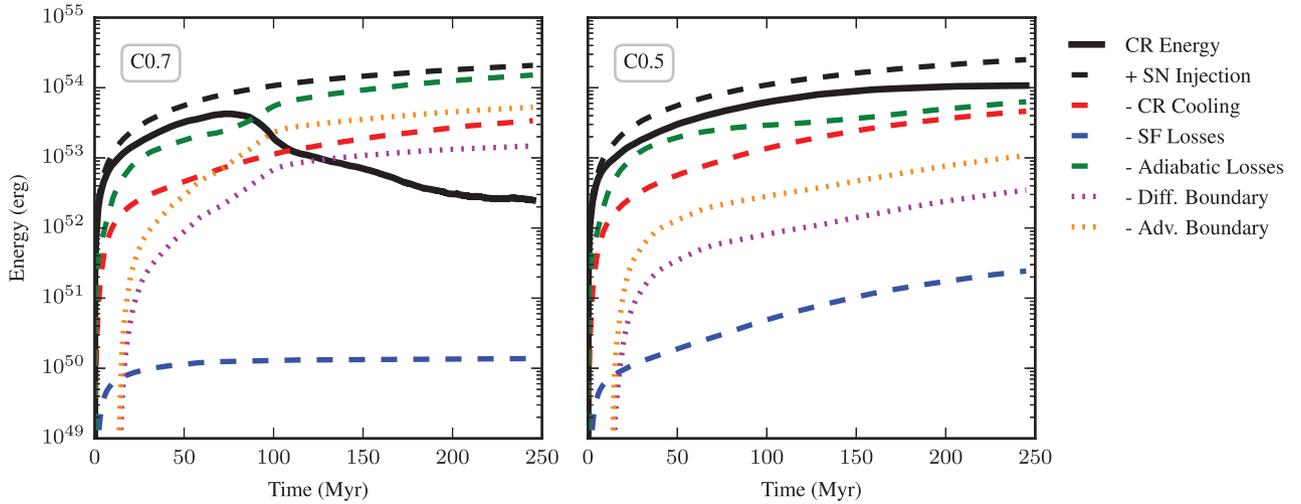
Figure 9. Comparison of the CR energy evolution for the C0.7 run (left-hand panel) which undergoes a thermal runaway and the C0.5 run (right-hand panel) which does not and has warm ($\sim 10^4$ K) gas filling most of the mid-plane volume. The quantities shown are summed over the entire simulation domain. The instantaneous CR energy is shown (solid black line) and source and loss quantities are shown as cumulative sums over time. The source term is the injection of CRs from SNe (black dashed). The loss terms shown are CR cooling from hadronic and Coulomb processes (red dashed); losses from the removal of star-forming gas from the mesh (blue dashed); adiabatic losses (green dashed); adiabatic losses over the outflow boundary (orange dotted); and diffusive losses from the outflow boundary cells (magenta dotted). The adiabatic boundary losses are the sum of the CR energy flux over the outflow boundaries from cells with faces on those boundaries. The diffusive boundary losses are the change in CR energy of boundary cells before and after the diffusion solver updates its solution.

other sink term relevant in the mid-plane, CR cooling, is not efficient enough to alter the trend in the CR energy evolution, and the total amount of CRs in the calculation stabilizes to a near equilibrium. The overall result is a medium where CR energy density dominates over thermal energy density.

3.4 Fast, slow, and no CRs: the impact of CR pressure and transport on the mid-plane structure

The phase structure of the gas impacts the evolution of the CR energy as we have shown, but how do the CRs feed back on these structures? How dynamically relevant are CRs to the phase structure evolution we have described?

We have already considered the impact of removing CRs from the calculation altogether. Fig. 4 gives a visual impression of the impact of CRs in the calculation, contrasting projections of gas density in runs with and without CRs. The models without CR diffusion do not track CR energy density and all SN energy is deposited in the form of thermal energy. The only energies in the gas are therefore thermal and kinetic. Figs 10 and 11 show a comparison of mid-plane gas properties and evolution between these sets of runs.

The FF models with and without CRs behave virtually identically in their volumetric mid-plane gas properties but show a difference in the overall height of the mid-plane gas layer and the global SFR. The difference in the mid-plane layer thickness is due to the enhancement of the overall pressure due to the CRs that puffs up the mid-plane gas. The enhanced SFR can be understood as the result of the effect of CRs on denser clouds which determine the global SFR in the FF model. The presence of CRs reduces the mass of dense gas at the mid-plane, which in turn depresses the SFR. We will further explore the impact of pressure on dense gas in Section 3.5.

The two selected MIX models behave differently with and without CRs. As shown in Fig. 10, the model with 70 per cent random SNe evolves more quickly to a thermal runaway without CRs: the mid-plane density drops more quickly and the mid-plane temperature rises more quickly. In the simulation with CRs, the median mid-

plane density does drop below $0.01 \, \text{cm}^{-3}$ and the temperature rises above 10^6 K, it just takes somewhat longer to do so. By contrast, in the run with 50 per cent random SN explosions, the two cases evolve to different states. The run without CRs undergoes a thermal runaway, but the run with CRs evolves to the peak driving state. We note that some of the differences seen in the FF model in the scale height and SFR are not at play in the MIX models in the same way. In the MIX models, the global SFR is purely a function of the mass of gas remaining in the volume (i.e. gas not lost due to outflows) and is not an indicator of the gas density.

Figs 4, 10, and 11 show that the fraction of random SN events necessary to precipitate a thermal runaway is higher and occurs over a much narrower range of f_{rand} when CRs are included.

In the setup tested here, the models with a 50 per cent random fraction represent a 'tipping point' of the dynamical evolution and are therefore an interesting case to consider. This value of f_{rand} represents a separatrix between two attractor solutions of the system: the thermal runaway and the self-regulated peak-driving solutions. CRs change where this tipping-point occurs in addition to softening the vertical gas density gradient and causing the emergence of cooler and smoother galactic outflows in comparison to purely SN driven, ballistic outflows (Simpson et al. 2016; Girichidis et al. 2018).

The microphysics underlying the choice of the CR diffusion coefficient κ is complex and depends on the specifics of CR plasma physics. A correct modelling of the plasma physics of CR transport would imply resolving the interplay of CR-driven instabilities (Kulsrud & Pearce 1969; Shalaby, Thomas & Pfrommer 2021) and various Alfvén wave damping processes (Guo & Oh 2008) so that the resulting level of Alfvén waves (in-)efficiently scatters CRs and thus set their transport speed (with a lower level of Alfvén waves implying less efficient CR scattering and thus faster diffusive transport). However, the large scale-separation between the CR gyroradius, where the wave-particle scattering happens, and the extent of galaxies precludes a direct plasma-dynamical modelling but instead argues for using a non-equilibrium fluid theory of CRs (Zweibel 2017; Jiang & Oh 2018; Thomas & Pfrommer 2019).
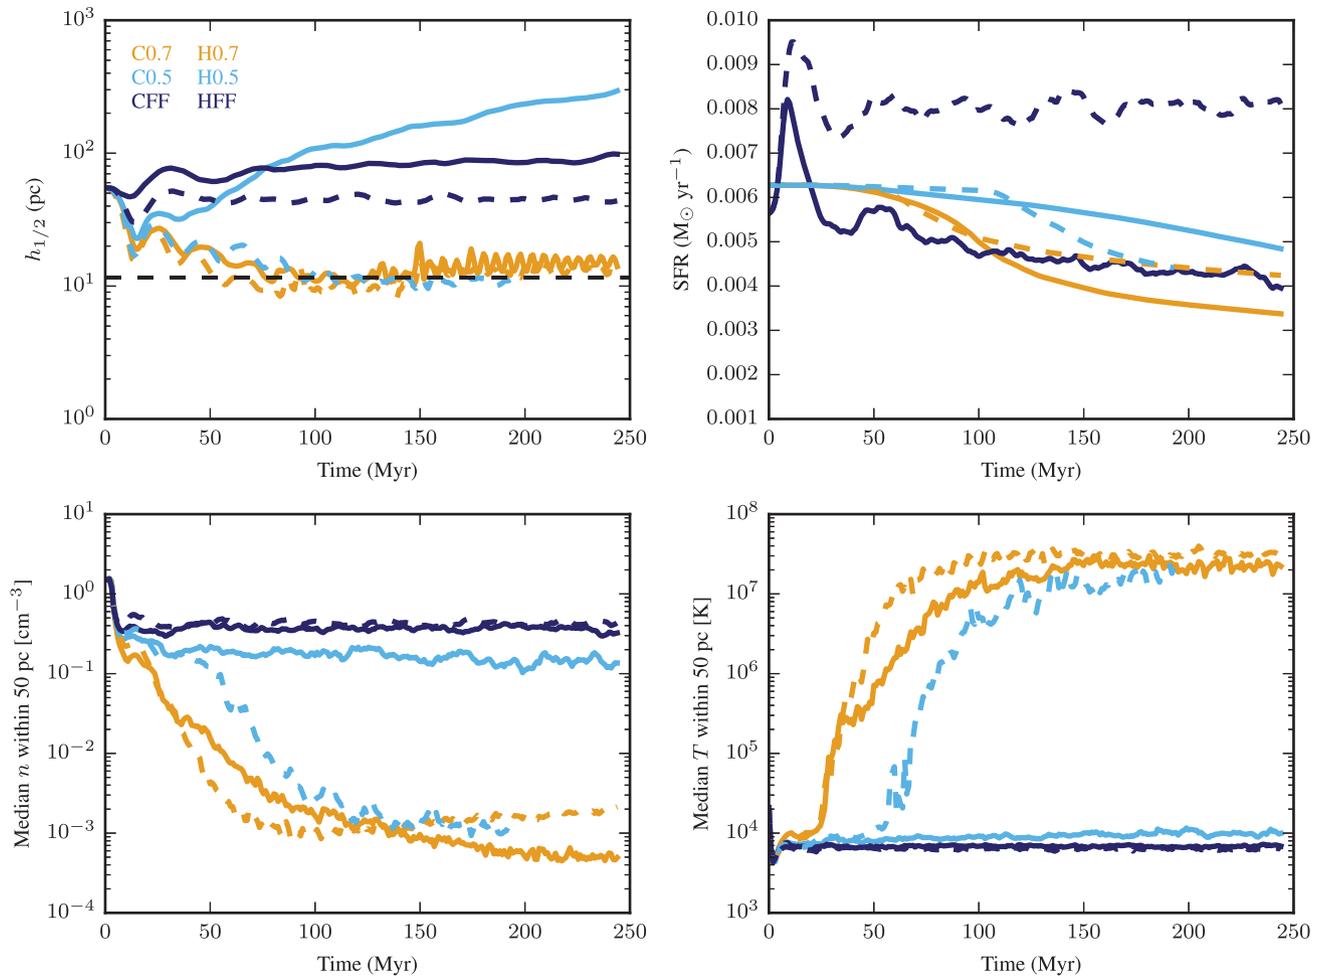
Figure 10. Evolution of models with (solid) and without (dashed) CRs compared. The panels shown are the same as Fig. 5. The SN explosion placement models shown are the FF model (pure peak-driving) and two MIX models with f_{rand} of 0.7 and 0.5. The $f_{rand} = 0.5$ models evolve differently with and without CRs. The presence of CRs is able to flip the model from the thermal runaway solution to the peak driving solution.

In the self-confinement picture of CR transport, the CR diffusion coefficient κ varies spatially and temporally as a result of the different balance of the Alfvén wave excitation and damping processes (Thomas & Pfrommer 2019; Thomas, Pfrommer & Pakmor 2021; Thomas & Pfrommer 2022) and CRs are transported as a superposition of advection, streaming and diffusion as can be inferred from the radio brightness profiles of radio-emitting non-thermal filaments in the Galactic centre (Thomas et al. 2020). Although this picture may be in tension with observations of CR spectra (Hopkins et al. 2022)

To study the sensitivity of our results to the diffusivity of the CR fluid, we have run the CFF and C0.5 models with different values of κ that bracket our fiducial value, between 10^{27} and 10^{29} cm² s⁻¹. This enables us to assess whether the locus of the separatrix discussed above also depends on the value of the CR diffusion coefficient, and therefore how sensitive our ISM model would be to the details of the self-confinement picture. Our fiducial value of κ is motivated by observations of the fluxes of secondary-to-primary nuclei, such as the boron-to-carbon ratio measured at the Earth (e.g. Evoli et al. 2020). We have also tested these models without diffusive transport ($\kappa = 0$), where the CR fluid is perfectly coupled to the thermal gas.

In our model, CRs diffuse along magnetic field lines and their diffusivity is set by the coefficient κ, which is the diffusion coefficient of the CR fluid in the direction of the local magnetic field according to equation (9) (in our model, the physical diffusion coefficient is 0

in all directions perpendicular to **b** so any remaining perpendicular diffusion is of a purely numerical nature). We have demonstrated how the presence of CR pressure in the mid-plane affects the fragmentation of the ISM. Varying the diffusivity of the CR fluid by adjusting κ will impact the energy density of CRs in the mid-plane, thereby influencing the ISM dynamics.

Fig. 12 shows the evolution of the mid-plane medium with different transport and SN explosion placement models. With the CFF model, runs with different κ evolve to the same mid-plane density structure, and show little difference from the non-CR models. There is a modest change in the mid-plane height of gas, where $h_{1/2}$ increases with decreasing κ.

With the C0.5 model, the mid-plane height of gas also increases with decreasing κ, but the effect is much more dramatic, over an order of magnitude. The mid-plane density evolution shows that while the fiducial model asymptotes to the peak driving solution, the model with a larger κ undergoes a thermal runaway, much like the model without CRs. The model with a lower κ undergoes a very different type of evolution. The model with low κ evacuates almost all of the gas from the mid-plane region because the CR pressure force acts for a longer time on the gas in comparison to the models with a larger κ, creating a mid-plane dominated by low density gas, but without the dense molecular clouds found in runs that undergo a thermal runaway.
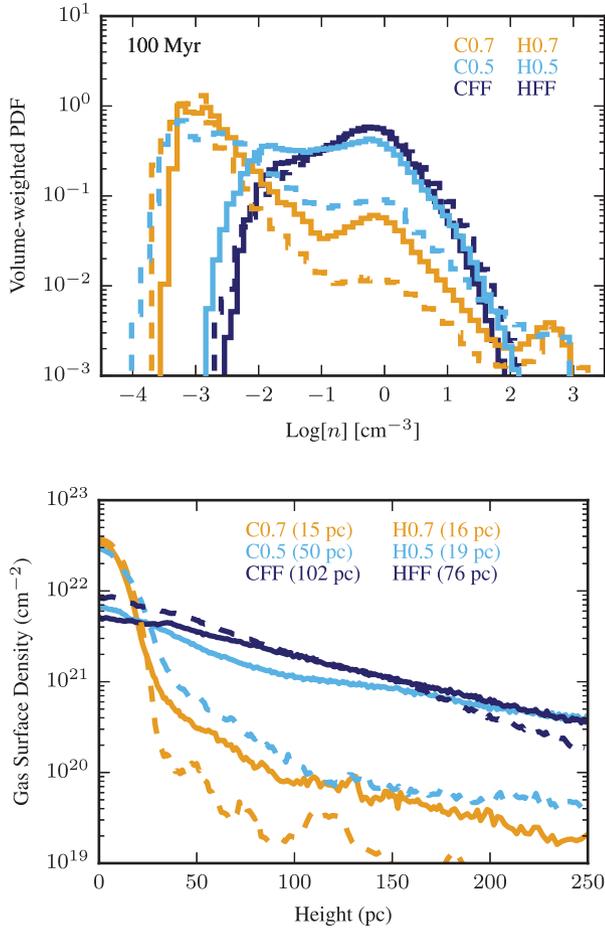
Figure 11. Comparison of simulations with CRs (solid lines) and without CRs or MHD (dashed lines) for models with the same SN explosion placement. The models shown are the same as in Fig. 10. The same quantities are shown as in Fig. 6: the distribution of gas densities within 50 pc of the mid-plane (top) and the vertical distribution of gas density (bottom). The presence of CRs changes the behaviour of models with $f_{rand} = 0.5$, but also thickens the amount of gas at heights about 100 pc in all models.

This is clear in Fig. 13 that shows the density and height distributions of mid-plane gas. In the CFF model there is a smooth progression with κ where the distribution of gas densities broadens with increasing κ, but the mid-plane layer thickens with more gas above 200 pc. In the C0.5 runs, faster diffusion also results in a thermal runaway and less gas at larger scale heights (although, more than in the pure hydro model). Slower diffusion prevents gas collapsing to densities higher than 1 cm^{-3} and leads to an inverted vertical gas density profile below 200 pc. The spatial variation in CR energy density in the high κ runs are quite smooth like the fiducial run seen in Fig. 7, whereas the lower κ runs show more variation as CR energy density takes longer to diffuse from injection sites. In terms of β_{CR}, the spatial variation in the high κ runs is similar to the fiducial model as well as in the low κ CFF run. In the the low κ C0.5 run, the evacuation of mass from the mid-plane creates a situation where the mid-plane thermal temperature is very high and therefore $\beta_{CR} > 1$ with pockets of $\beta_{CR} < 1$ where some denser gas survives.

CR models with no diffusive transport (where the CR energy density is perfectly coupled to the thermal gas) produce a thickened gas layer. In C0.5-adv, the formation of dense clouds is suppressed as in the low diffusion case. However, the density distribution is skewed to higher densities in the no diffusion case, resulting in a medium

that is overall denser than the low diffusion case. This is likely due to the lack of any diffusive transport, without which, gas remains confined to the mid-plane, albeit in a layer that is very puffy with a slight inversion in the central density profile.

From this analysis, we can conclude in the model tested here that the diffusivity of the CRs has more of an impact on the larger scale structure of the ISM (e.g. the scale height of the gas), and less on the mid-plane structure, unless the value of κ is extremely high or low. High values of κ result in a model where CRs become unimportant in the mid-plane, and the evolution asymptotes to they model without CRs. Low values of κ can result in the suppression of dense clouds, except in the case where CR energy is continuously deposited in dense gas and CR cooling acts to remove CRs.

3.5 Effective pressure support

As we have seen, the inclusion of CRs can impede or slow the fragmentation of gas into different phases. In models where the diffusive transport of CRs is slowed, the distribution of gas densities narrows.

To better understand this behaviour, we consider the effective pressure and how it influences the gas evolution. The effective pressure is defined as

$$P_{\text{eff}} = P_{\text{therm}} + P_{\text{CR}} = (\gamma - 1)\varepsilon + (\gamma_{\text{CR}} - 1)\varepsilon_{\text{CR}}. \quad (21)$$

This definition of pressure combines thermal pressure and CR pressure into a single measure of the pressure. We neglect magnetic pressure because we found that it is almost always subdominant (see Fig. 7).

Fig. 14 shows the evolution of the distribution of effective pressure in our fiducial CR runs within 50 pc of the mid-plane (unlike Fig. 1 which shows the thermal pressure in a larger volume).

This gross behaviour can be understood by considering the pressure support in the dense gas. As gas cools, dense regions contract to maintain a constant thermal pressure (i.e. isobaric cooling). In models with a sufficiently large reservoir of CRs, the CRs become important. As dense gas contracts, CRs become adiabatically compressed which produces an outward CR pressure gradient force that if steep enough can stabilize the gas against collapse. Therefore, in the presence of CRs, it is more difficult for gas to become denser from radiative cooling.

This can be seen in Fig. 14. The CRs provide a minimum effective pressure support in the gas that is constant across density (seen in C0.5 and CFF). When gas radiatively cools in these models, pressure forces do not act to make gas more dense, because the CR pressure maintains an equilibrium.

However, CRs cool and so need to be replenished to provide effective pressure support. This occurs in the peak-driving case, where CRs are replenished by direct injection via SN explosions in dense gas. In the thermal runaway case (e.g. C0.7 shown in Fig. 14), the only potentially efficient source of CR replenishment is diffusion of CRs from lower density gas into clouds. But for our diffusivity, this process is not fast enough to overcome the radiative cooling. The dense gas will therefore collapse further. Ultimately, only our imposed Riemann pressure floor will halt collapse.

4 OBSERVATIONAL SIGNATURES

Here we explore predictions for the emission of gamma-rays arising from interactions between the CR fluid and the thermal gas and the distribution of H I atoms that give rise to 21 cm emission.
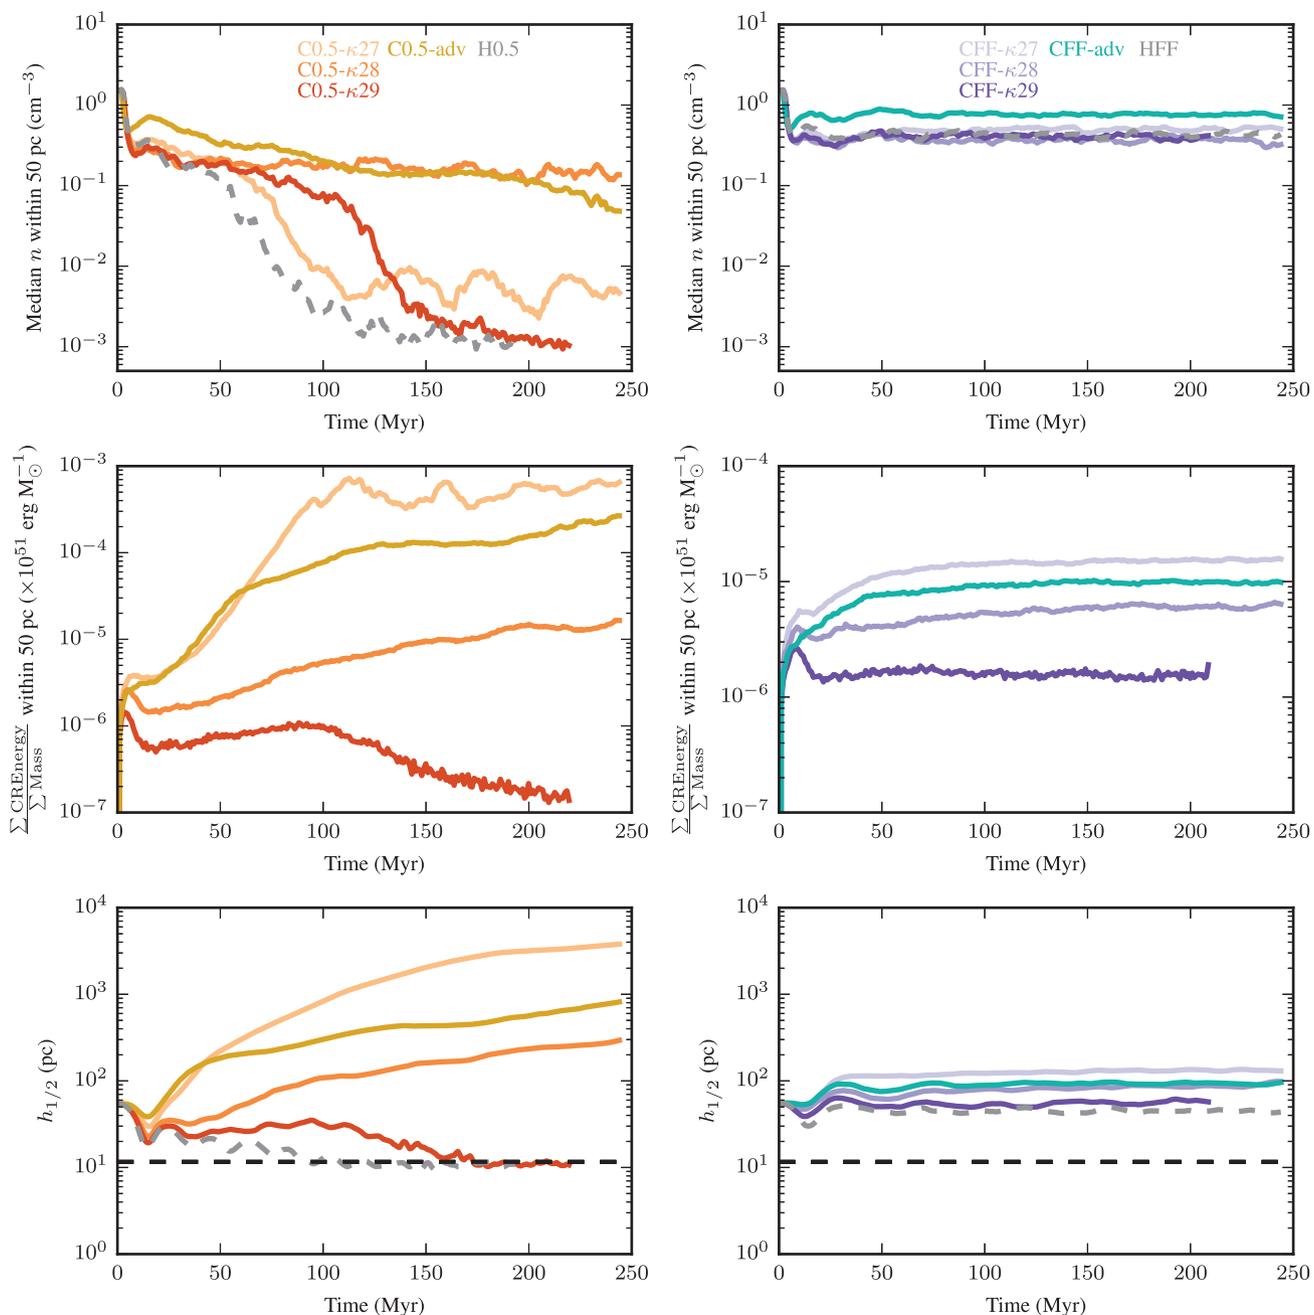
Figure 12. Evolution of simulations run with varying κ, the global anisotropic diffusivity of the CR fluid. Top row: the median number density of the thermal gas within 50 pc of the mid-plane over time. Middle panel: the total CR energy within 50 pc of the mid-plane divided by the total gas mass in the same volume over time. Bottom panel: the height containing half the original gas mass of the simulations. Two models for SNe placement are shown, the MIX0.5 model (left-hand panel) and the FF model (right-hand panel). For both models the hydrodynamics-only model are shown with a dashed line. Models with the CR fluid but evolved without diffusive transport are also shown. As κ decreases, the scale height of the mid-plane gas layer increases, but the phase structure of the mid-plane gas is still sensitive to the SN explosion placement model.

4.1 Hadronic gamma-ray emission

We follow Pfrommer et al. (2017b) in modelling gamma-ray emission from interactions of CRs with the ISM. This model assumes gamma-rays come from two sources: the primary channel from the decay of neutral pions created by inelastic collisions of CRs with thermal nucleons and the secondary channel from inverse Compton scattering of relativistic electrons. These processes are dependent on the energy spectrum of CRs, which we do not self-consistently simulate, but instead we assume a power-law CR momentum spectrum. We focus

on the primary production channel of gamma-rays through pion decay, which is dominant for at the star formation rates in galaxies the mass of the Milky-Way and greater (Werhahn et al. 2021b). The gamma-ray emission from a cell (luminosity per unit volume) in an energy interval $[E_1, E_2]$ comes from integrating the source function $s_{\pi^0-\gamma}(E_\gamma)$ of decaying neutral pions (adopting the semi-analytical model of Pfrommer & Enßlin 2004):

$$\Lambda_{\pi^0-\gamma} = \int_{E_1}^{E_2} s_{\pi^0-\gamma}(E)E\,\mathrm{d}E \tag{22}$$
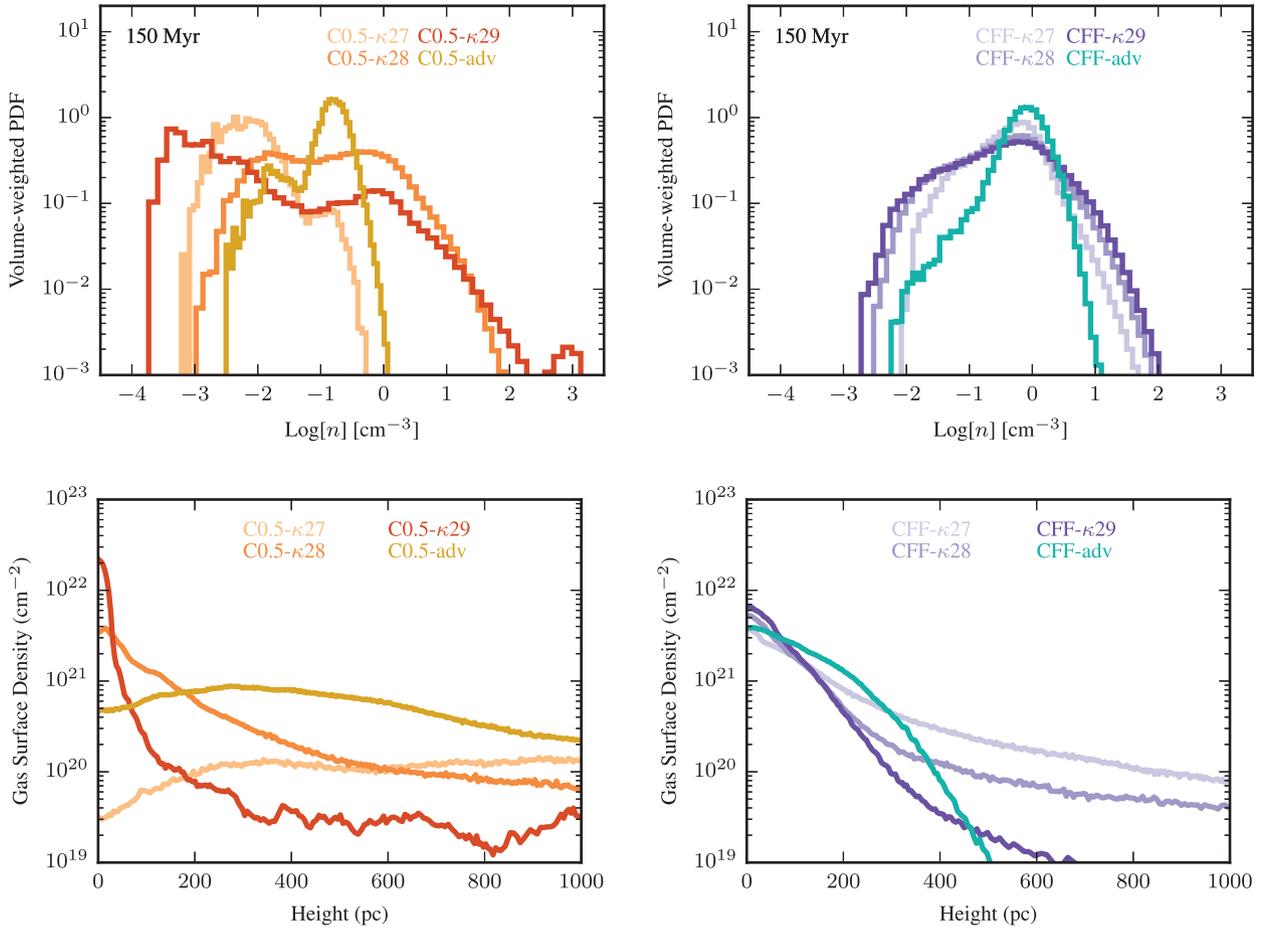
Figure 13. Properties of models with different CR diffusivities after 150 Myr of evolution, including CR models with no diffusion. Two SNe placement models are shown: one with SNe placed only in high density peaks (CFF, left-hand panel) and one where 50 per cent of explosions are positioned randomly in the mid-plane (C0.5, right-hand panel). The fiducial CR model is run with $\kappa = 10^{28}$ cm^{2}s^{-1}. Top: Volume-weighted probability density functions (PDFs) of the logarithm of the gas number density within 50 pc of the mid-plane. Bottom: Surface number density of gas along the vertical direction. The simulation volumes above and below the mid-plane are combined and the mid-plane is at height 0.

We compute the luminosity for each cell in the Fermi energy band (0.1–100 GeV).

We use an estimate for the gamma-ray luminosity from Pfrommer et al. ([2017b](#)) who assume a power-law distribution for the momentum of CRs:

$$f(p) = C_p p^{-\alpha} \theta(p - q), \tag{23}$$

where $p = P_p/m_p c$ is a dimension-less proton momentum (m_p is the proton mass and c is the speed of light). The function θ is the Heavyside step function and q is the low-threshold momentum limit. We adopt $q = 0.5$ to account for the fact that low-energy CRs are efficiently cooling via Coulomb and ionization interactions (Werhahn et al. [2021a](#), [b](#)). We choose a fixed power-law slope of $\alpha = 2.05$. A more natural choice for the power-law slope would seem to be $\alpha = 2.2$ (Caprioli et al. [2020](#); Haggerty & Caprioli [2020](#)), corresponding also to the slope that maximizes pion decay emission in the Fermi band. However, this slope, as applied in this analytic model, has been shown to be effectively degenerate with the choice of CR SN injection efficiency (Pfrommer et al. [2017b](#)). Our chosen injection efficiency of 10 per cent, while within acceptable limits, has recently been shown to likely be an overestimate, as SN acceleration models that account for the magnetic field-shock front obliquity favour a fraction closer to 5 per cent when integrated over the entire remnant, and independent

of the magnetic coherence length (Caprioli & Spitkovsky [2014](#); Pais et al. [2018](#)). To account for this, we have chosen the value of α appropriate for this factor of two higher injection efficiency, i.e. one that results in a factor of two lower pion decay emission.

The normalization factor C_p is determined for each cell by equating the simulated energy density of CRs to the integrated energy density from our assumed $f(p)$ (which is a good approximation for the CR spectrum above the kinematic threshold $p = 0.78$ of the hadronic reaction; see Girichidis et al. ([2022](#)) for a simulated CR spectrum in a galaxy simulation):

$$\varepsilon_{CR} = \int_0^\infty f(p)E(p)\mathrm{d}p, \tag{24}$$

where $E(p) = (\sqrt{1 + p^2} - 1)m_p c^2$ is the kinetic proton energy. Werhahn et al. ([2021b](#)) demonstrate (in their Appendix B) that this analytical model reproduces the result of a more detailed numerical integration that includes the energy-dependent hadronic cross section with an accuracy of 10 per cent.

Under these assumptions, the source function and hence the luminosity density of gamma-rays are proportional to the energy density of CRs times the number density of target nucleons for the hadronic reaction: $\Lambda_{\pi^0 - \gamma} \propto \varepsilon_{CR} n_n$. The spatial distribution of gamma rays is therefore highly dependent on the density structure of
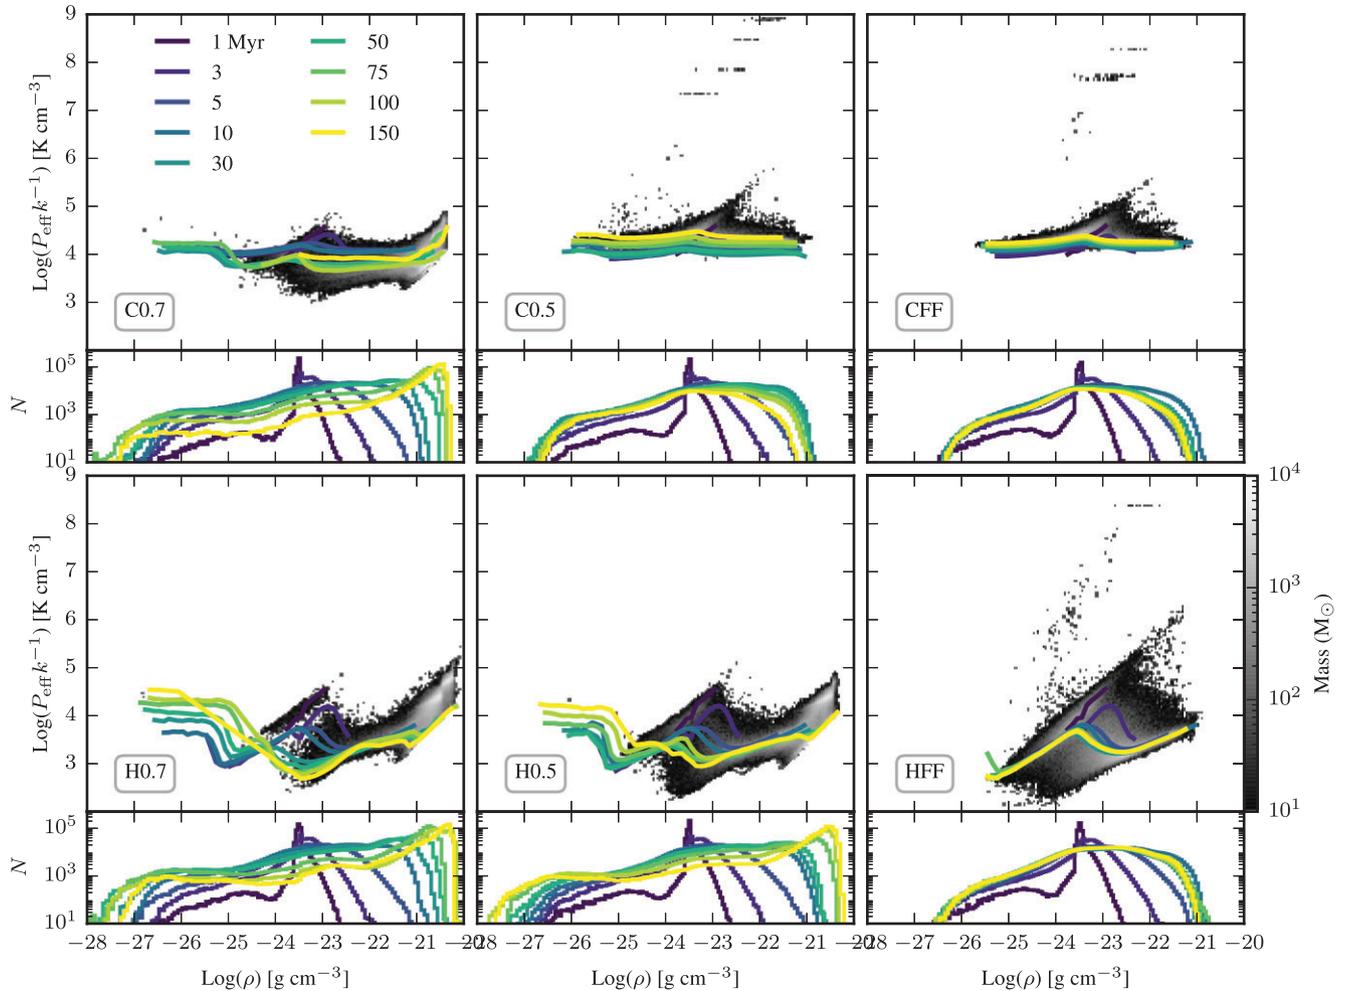
Figure 14. The effective pressure evolution for three different SN explosion placement models run with (top row) and without (bottom row) CRs is shown. Panels show the effective pressure versus gas density after 100 Myr of evolution within 50 pc of the mid-plane in greyscale. For the times indicated, curves following the median trend in this space are also shown. Below the pressure panels, the distribution of cell gas densities are shown.

the thermal gas (because $n_n = \rho/m_p$) and the speed with which CRs are transported from their birth sites.

The total gamma-ray luminosity in the Fermi band from pion decay in the patch of ISM simulated here can be computed by summing the emission from all cells in the simulation volume. The evolution of this luminosity is shown in Fig. 15, which shows that the time variability of gamma-ray emission differs between the two solutions discussed in Section 3. In the thermal runaway models, the emission is highly variable. It is especially variable in models with a non-zero number of SNe in high density peaks, like C0.7 and C0.6. The models that asymptote to the peak driving solution have very minimal time variability. The overall level of emission is very similar between all models except for the purely random placement model, C1.0. That model has an overall luminosity an order of magnitude lower than the other models.

The reasons for these differences is revealed in Fig. 7. This figure shows that fast diffusion in the mid-plane creates a roughly spatially uniform distribution of CR energy throughout the mid-plane. The level and variability of gamma-ray emission is therefore set by the distribution and volume-filling factor of nucleons. In peak-driving models where the mid-plane volume is mostly filled by warm neutral gas and where there are smaller gas density perturbations, the emission from recent SNe is not very different

from the background gamma-ray emission coming from this warm gas.

In the thermal-runaway models, most of the mid-plane volume is filled by very low density ionized gas, and most nucleons are found in the molecular clouds that occupy a very small volume of the mid-plane. The background gamma-ray emission is therefore small compared to the emission that arises when SNe occur near dense clouds. The variability of the thermal-runaway models are therefore a result the concentration of target nucleons in molecular clouds and the MIX SNe placement model. Periods of high emission arise when a recently injected SNe coincides spatially with a molecular cloud and low emission when randomly injected events dominate. Freshly injected CR energy near a molecular cloud does not remain spatially concentrated for long since we have no model for variable diffusion (such as Semenov, Kravtsov & Caprioli 2021). The energy quickly diffuses, causing the emission of gamma-rays to drop in the region of the molecular cloud and therefore the entire box. In C1.0 where all SNe are injected randomly, the coincidence of a recent SNe and a molecular cloud is very rare (by construction) so C1.0 has less variability and an overall lower level of emission. The runs that have the highest degree of variability are runs where there is a large f_{rand}, but with some non-zero seeding of SNe in dense gas where the placement of SNe in molecular clouds light them up in gamma-rays.
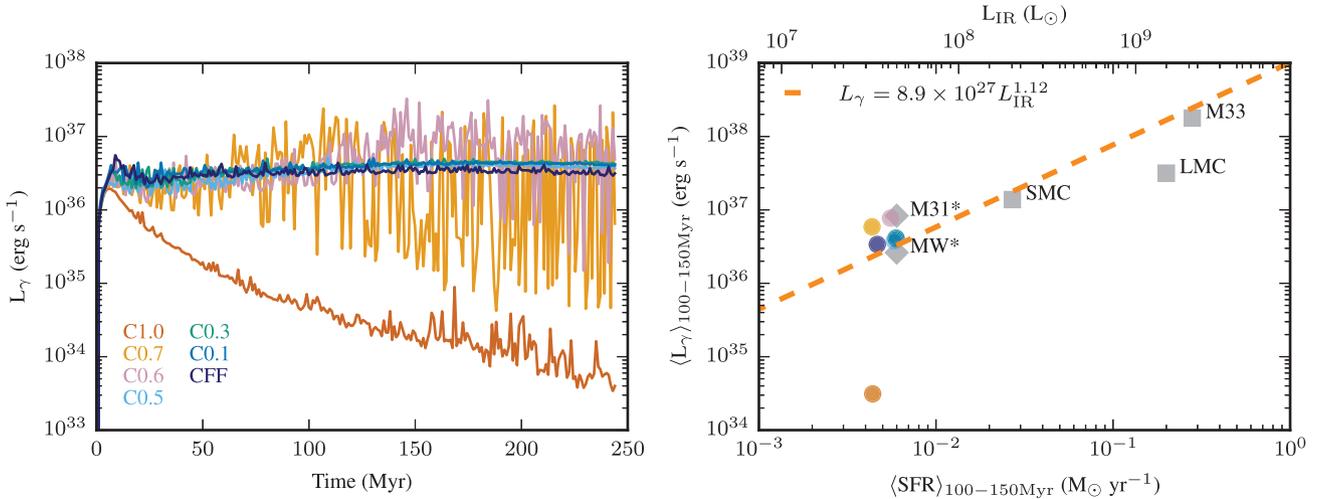
Figure 15. Gamma-ray emission arising from pion decay in the Fermi band of 0.1 to 100 GeV. Left-hand panel: Time evolution of the total gamma ray luminosity of each simulation volume. Thermal-runaway models demonstrate a higher degree of variability than peak driving runs. Right-hand panel: The average gamma-ray luminosity averaged between 100 and 150 Myr for all the runs. An extrapolation of the best-fitting line to galaxies observed both with Fermi and in the infrared is shown (dashed orange line) assuming a relation between infrared emission and SFR from Kennicutt. Observed values for the SMC, LMC, and M33 are shown (squares) and values for the MW and M31 are shown scaled to an SFR of 6×10^{-3} $M_\odot yr^{-1}$ (diamonds), corresponding to the local patch of ISM simulated here.
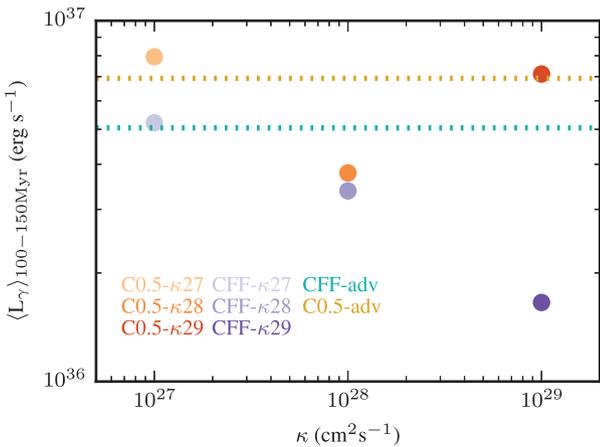
Figure 16. Gamma-ray emission arising from pion decay in the Fermi band of 0.1 to 100 GeV averaged between 100 and 150 Myr. Models run with different CR diffusivities κ are shown with both the CFF model (purple points) and the C0.5 model (orange points). Models run with CRs but no diffusive transport (advection only models) are shown as dotted lines for the CFF model and C0.5 model.

The relation between gamma-ray luminosity and SFR is also seen in Fig. 15. Galactic-scale gamma-ray emission is observed to correlate with FIR emission over several orders of magnitude for systems where Fermi has been able to make measurements, from the SMC to Arp 220 (Ackermann et al. 2012; Rojas-Bravo & Araya 2016). Our models do not probe a range of star-forming environments so we can only probe this relation at a single SFR; however, we see that the average level of gamma-ray emission does lie on the observed galactic-scale trend. The one exception is C1.0, which as we have explained has a lower level of gamma-ray emission.

Fig. 16 shows how the gamma-ray luminosity varies with the CR diffusion coefficient, κ. The luminosity can vary by a factor of several with κ, with lower κ models trending toward higher luminosities, especially for the CFF models, thus matching the pure CR advection

models with $\kappa = 0$. This is due to the accumulation of CR energy in the mid-plane that results from the slower diffusion. It is also the case, however, that the presence of CRs impacts the amount of dense gas in the mid-plane, the gas that is most luminous in gamma-rays. In C0.5-κ29, the fast diffusion of CRs reduces the amount of CRs relative to mass in the mid-plane (see Fig. 12), however, the reduced CR pressure allows more dense gas to form (see Fig. 13). Even with the reduced CR energy density in the mid-plane, the enhanced dense gas means that the model has on average an enhanced gamma-ray luminosity.

4.2 Atomic hydrogen

Fig. 17 shows the vertical column density of H I and H$_2$ gas in models with and without CRs. The scale height, temperature, and composition of neutral gas varies between injection models, but also changes with and without CRs.

Some of these differences are due primarily to the differing solutions to which these models evolve, but some differences seem to be due to the presence of CRs irrespective of the mid-plane's ISM structure. Models that evolve to the peak-driving solution have more vertically extended H I and lower surface densities of H$_2$ when CRs are included. Thermal runaway models have less H I gas above the mid-plane and very concentrated H$_2$ layers. The peak surface density of H$_2$ in all the thermal runaway models saturates at a value of $\sim 3 \times 10^{22}$ cm^{-2}. This constant value is likely due to the Riemann pressure floor imposed in dense gas.

Part of the extended H I gas in the CR runs may be part of the base of a CR-driven outflow, which we will explore in future work. The overall thickness of the H I layer varies between approximately 20 pc (C1.0) and 150 pc (C0.1) in these runs while it is typically less than 50 pc in the no CR runs. We define these heights as the 1 e-folding height, i.e. where the gas surface density drops by 1/e.

In all models, the central mid-plane H I gas is dominated by cold gas with temperatures less than 10^3 K. The fraction of H I that is composed of this cold neutral medium (CNM) varies from 60 to
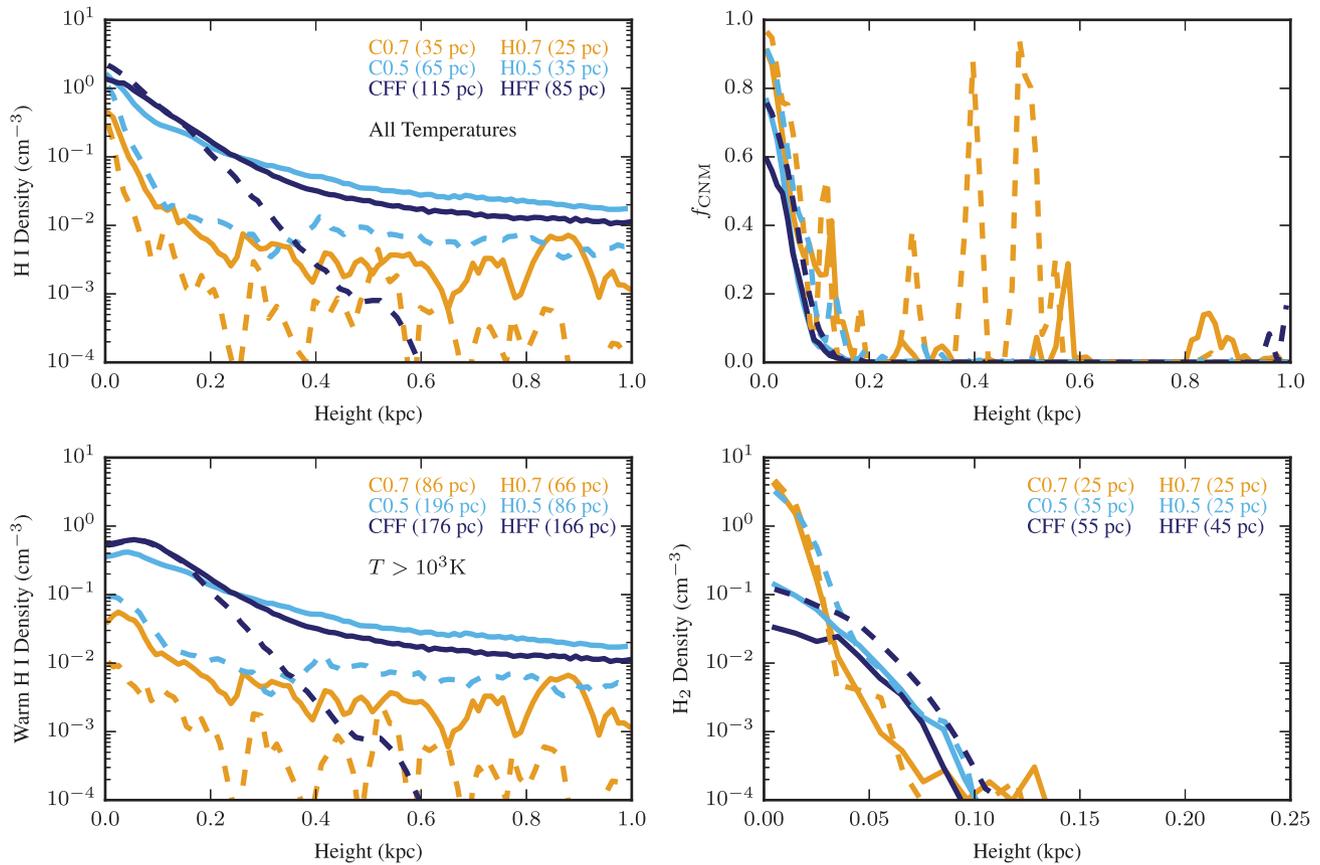
Figure 17. The vertical distributions of different species of neutral hydrogen shown after 100 Myr of evolution for simulations with CRs (solid lines) and without CRs (dashed lines). Top left-hand panel: The total number density of H I atoms versus height. The e-folding scale heights of the density are indicated. All gas regardless of temperature is included in this plot. Top right-hand panel: The fraction of H I atoms in cold gas ($T < 10^3$ K), representing the mass fraction of the CNM. Bottom left-hand panel: The number density of H I atoms in gas that is warm ($T > 10^3$ K). This gas corresponds to the optically thin WNM. Bottom right-hand panel: The number density of H_2 molecules.

nearly 100 per cent as shown in Fig. 17. Even at a height of 100 pc, the CNM fraction (f_{CNM}) can vary between 20 and 40 per cent. These high fractions indicate that a substantial amount of H I gas may be optically thick in the mid-plane, which would impact the column density of H I emission at 21 cm.

The height of the MW's 21 cm emission at the solar circle is approximately 150 pc as defined as the half-width half-maximum (Kalberla & Dedes 2008), larger than the upper range of our models. However, given the large central CNM fractions in our simulations, even up to a height of 100 pc, the scale height of the H I volume density should be regarded as a lower limit for any inferred height of the 21 cm emission because a large portion of the cold H I will be optically thick. We can consider the warm H I, with temperatures greater than 10^3 K, as providing an upper limit to the scale height of the 21 cm emission, as this warm H I will certainly be optically thin. Fig. 17 shows that the e-folding heights of the warm H I atoms are substantially larger than the total H I. The height of the gas layer of this optically thin H I is typically greater than 150 pc for peak-driving runs. For thermal runaway models, the WNM is also more extended, but does not approach 150 pc in scale height.
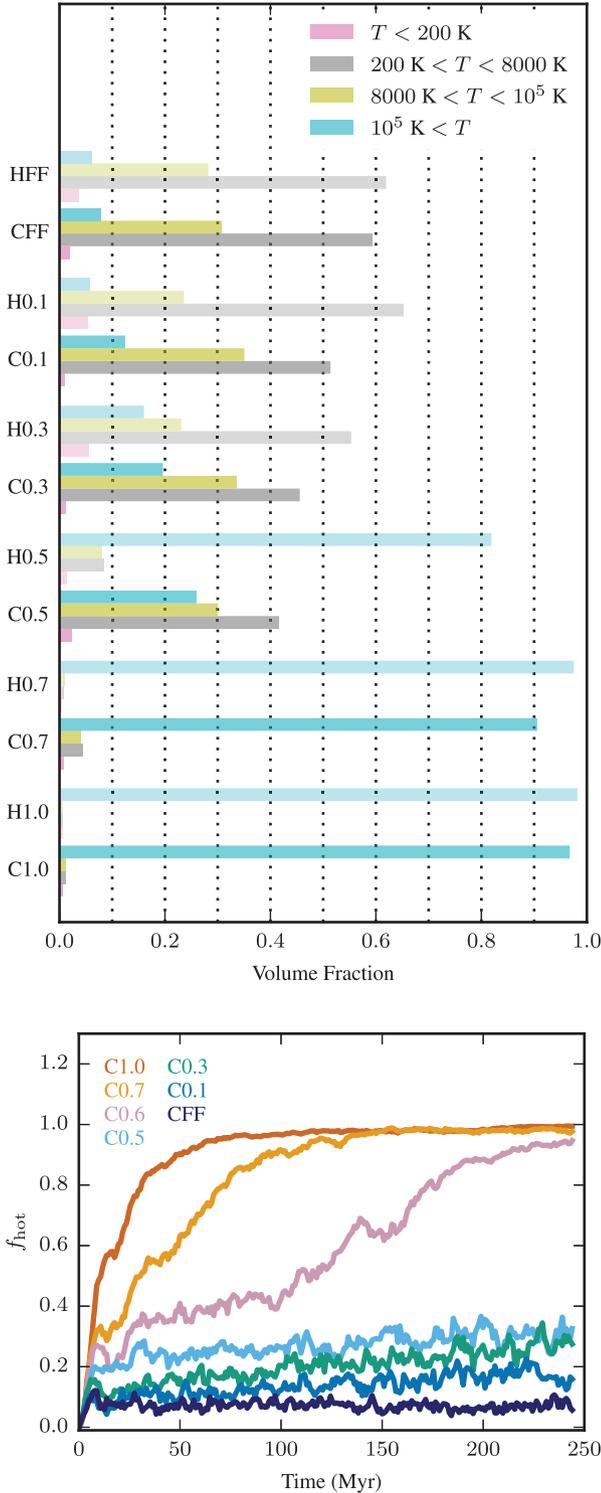
We also present in Fig. 17 the distribution of H_2. It should be noted, however, in order to capture a converged H_2 component, much higher resolution is needed (Joshi et al. 2019). Zoom simulations of molecular clouds suggest that at our resolution, the underestimate of the molecular gas fraction is of the order of a factor of two (Seifried et al. 2017).

Our models have demonstrated that CRs enhance the thickness of the H I layer. With the set of physics we consider in this study, CRs appear to thicken the H I layer to a degree that brings models, with varying SN placement assumptions, in better agreement with the MW's 21 cm emission.

5 DISCUSSION

The ISM system examined here includes many of the physical effects found in the real ISM of the MW, but there are also many effects we have not considered in this study. Comparing the solutions we have identified here in this model to the state of the MW's ISM can give us some insights into this other physics.

Fig. 18 shows the fraction of the mid-plane volume occupied by gas of different temperatures. The transition between the peak driving and thermal runaway solutions is evident in these fractions. In the MW, measurements of gas in these temperature ranges are dependent on different observables that are sensitive to density and ionization state, in addition to temperature. Generally, the hot ionized medium (HIM) is thought to fill approximately 50 per cent of the volume. The warm ionized medium (WIM), gas of roughly 10^4 K found in H II regions, comprises about 20–30 per cent of the volume, and the warm neutral medium (WNM), gas of several thousand K and traced by warm H I gas, comprises about 30–40 per cent. Cold gas, with temperatures ≤ 100 K and traced by cold H I (the CNM) and

that 30 per cent of the volume). Other studies have found that a set of clustered SN explosions or SN explosions from dynamically formed sink particles are able to match the volume fraction of the HIM (e.g. Rathjen et al. 2021). We find that as the models in the peak driving phase get closer to tipping over to a thermal runaway, the volume fraction occupied by the HIM can pass through 50 per cent, but it is not a stable state with the physical model tested here.

The WNM and WIM are completely suppressed in simulations with a thermal runaway, unlike the MW where combined, they comprise about half the ISM volume. In the peak driving models, they occupy an appreciable amount of the volume as seen in Fig. 18. We separate the gas into temperature ranges here only, but we see that the cooler gas component (grey bands with $200\,K < T < 8000\,K$) is the dominant component of the medium, typically 40 per cent of the volume when $f_{rand} = 0.5$ and trending to 60 per cent in the FF models. The component at higher temperatures ($8000\,K < T < 10^5\,K$) seems to be more consistent in its volume fraction across peak driving runs (~25 per cent). This gas appears to be cooling SN remnants placed in the dense gas. They are constantly replenished by the injection model as they constantly cool away due to thermal radiative cooling.

This origin for the WIM is different from the MW where it is typically associated with H II regions produced by massive stars. While the simulations presented here have an interstellar radiation field attenuated by dense gas, they do not include a dynamic radiation field that includes radiation transport. Kannan et al. (2020) in models similar to those presented here tested the effects of photoionzation and radiation pressure with full radiation transport and found an enhancement in the supply of warm gas ($10^3\,K < T < 10^5\,K$). The combination of radiation and CR transport will likely combine in a non-linear fashion which we will explore in future work. However, the simulations of Rathjen et al. (2021) that include both effects suggest that the inclusion of radiation is especially important for capturing the warm phase.

Beyond radiation transport, the models here neglect by construction many effects of disc dynamics and the cosmological environment of the ISM. Turbulence plays an important role in the development of dense gas and in our model is seeded primarily by SN explosions, however, disc models suggest that gravity driven turbulence may be more consistent with observations of gas velocity dispersions (Krumholz & Burkhart 2016). Disc shearing likely plays an important role in the creation of dense gas, even at the resolution considered here (e.g. Smith et al. 2014). The stratified gas layer considered here allows for the outflow of gas from the mid-plane, but there is no CGM included that would be a potential reservoir of gas inflow to the disc or a mediator for fountain flows that arise from the disc but fall back (Tumlinson, Peeples & Werk 2017). Gas inflow to the Milky Way has been estimated at approximately $10^{-3}\,M_\odot kpc^{-2}\,yr^{-1}$ (Fox et al. 2019; Werk et al. 2019), approximately 10–20 per cent of the gas outflow rates found in these models. However, we note that without the physics associated with cloud disruption (local UV heating from massive stars, shearing forces, and the creation of star particles that remove gas) gas injected into the ISM would likely be locked up into dense clouds and not promote cycling of the ISM phases. What may happen is gas inflows from the CGM may affect the dynamics of outflows launched from the disc, which would possibly affect the gas cycling at the halo–disc interface.

While these models reach a high resolution of several pc in the mid-plane (see Fig. 2) and this is likely sufficient to capture the cooling of SN energy, we do not have the resolution to capture the development of SN remnant bubbles self-consistently and in the development of the top layers of the mid-plane, the steep gradient in gas density produces a steep gradient in cell sizes (although limited by our adjacent volume limit refinement criterion). Numerical mixing

Figure 18. Top: Fraction of the mid-plane volume (within 50 pc of the mid-plane) occupied by gas in different temperature ranges. Simulations with (dark colours) and without (light colours) CRs are compared with varying SN explosion placement models. These volume fractions are averaged between 95 and 104 Myr. Bottom: The evolution of the volume fraction of gas with temperatures greater than 10^5 K within 50 pc of the mid-plane.

molecules, is of order one per cent of the volume (e.g. Cox 2005; Draine 2011).

In our model, the HIM is either too volume-filling (occupying over 80 per cent of the volume) or not volume-filling enough (less

in these top layers may result in a suppression of dense structures or different turbulent properties.

We have presented a description of the range of behaviours for the ISM in the presence of a diffusive CR fluid with a constant κ. Models of SN remnants suggest that non-resonant streaming instabilities (e.g. Bell 2004) may lead to a decrease in CR diffusion away from SN remnants (Schroer et al. 2022). Semenov et al. (2021) test on much larger scales the impact of spatially varying CR diffusion in galactic discs and find that suppressing the diffusion of CRs away from birth sites enhances disc stability and prevents the fragmentation of gas. In addition, the interplay of resonant CR streaming instabilities (Kulsrud & Pearce 1969; Shalaby et al. 2021) and various wave damping processes should establish an effective CR propagation speed that varies spatially and temporally. Thomas, Pfrommer & Pakmor (2022) found that while the CR diffusion coefficient reaches a steady state in most galactic environments, CR transport itself appears not to reach a steady state and therefore, cannot be well described by either CR streaming, CR diffusion or a combination of both paradigms but instead obeys non-equilibrium transport mediated by resonant Alfvén waves.

The simple effect of Alfvén wave cooling of CR energy may play a role (Wiener, Pfrommer & Oh 2017) and is neglected here. We found losses from hadronic and Coulomb cooling to be subdominant to adiabatic losses, however Alfvén cooling has been demonstrated to have an effect on the temperature and structure of CGM gas (Buck et al. 2020). If the losses of CR energy are significant in the ISM, the effects of CR pressure on the phase structure and evolution of the ISM would be reduced. Across the models we have tested, in situations where CR energy density becomes less dominant, models tend towards the thermal runaway case (e.g. Fig. 8).

These plasma effects may shift the transition $f_{\rm rand}$ in our model between the peak driving and thermal runaway solutions. This is because confining CRs injected in dense gas within clouds longer would require a less frequent replenishment of fresh CRs to alter the cloud's collapse. However, if κ is allowed to increase away from the sites of recent SN explosions, CRs may escape the medium so quickly that effects like the enhanced H I layer thickness may not materialize. We note that while we do not vary κ, the CR diffusion time does vary with gas density as shown in Fig. A3. It is actually fastest in high density gas, which is likely the opposite trend that would result from a smaller κ near sites of SN explosions coincident with dense gas.

6 CONCLUSIONS

The models we have described have allowed us to test the behaviour of a diffusive CR fluid under specific astrophysical conditions, i.e. the application of SN explosions to a magnetized ISM at a prescribed rate in following a prescribed spatial distribution. We have varied these astrophysical variables to test the tolerance of the collisional and diffusive fluids that comprise our simulated ISM system. We have neglected many important effects for the evolution of the ISM: clustered SN explosions, disc dynamics, spatially and time varying radiation transport, and ISM–CGM interactions. This was intentionally adopted to provide a simpler system with which to focus on the unadulterated impact of CRs at high resolution where we are capturing resolved SN explosions. This approach lays a foundation on which we can better interpret a richer set of astrophysical effects in future work.

Our main results are as follows:

(i) We have systematically varied the amount of SN energy injected into high-density gas versus random locations, and consistent with previous works, models with large amounts of SN energy injected into dense gas evolve to a 'peak-driving' solution where warm gas is volume filling. Models with more SN energy injected into random locations develop a 'thermal runaway' where hot gas becomes volume filling.

(ii) The transition between these two solutions is mediated by CRs. As seen visually in Fig. 4, when CRs are present, the medium only fragments when more SN energy is coupled to random locations (i.e. into low density gas). We see in Figs 10 and 11 that CRs can entirely prevent the development of the thermal runaway over the 200 Myr run time of our simulations with some placement models.

(iii) We examine the partition of energy within the mid-plane of our simulations and find on global scales, turbulent and thermal energies are typically greater than CR energy when the medium is dominated by $T > 10^6$ K gas and less when it is dominated by warm gas ($T < 10^4$ K) as shown in Fig. 8. However there can be a great deal of spatial variation on pc scales in the distribution of thermal, kinetic, and magnetic energies while fast diffusion of the CR fluid means it can maintain an approximately spatially uniform distribution within the mid-plane (see Fig. 7)

(iv) The evolution of CR energy within simulations is a balance between injection through the SN model and a variety of sink terms, the dominant one being adiabatic losses from the expanding mid-plane gas (see Fig. 9). Hadronic and Couloumb losses can be significant as well especially in simulations dominated by warm gas.

(v) We tested different values of κ, the diffusivity of the CR fluid, in the presence of two different SN placement models. The comparison of these simulations is shown in Figs 12 and 13. We find that fast diffusion results in models similar to the hydrodynamics only models. Slow diffusion results in the build up of CR energy in the mid-plane. The result of this build up can vary depending on the SN placement model, with some models creating pressure gradients so strong that gas is evacuated from the mid-plane.

(vi) We examined the evolution of the effective pressure in the mid-plane and show that CRs impose a minimum pressure floor in the medium across densities (Fig. 14). In cases where evolution of the thermal pressure is significant, this pressure floor can break, allowing for the fragmentation of the medium due to SN explosions. Otherwise, the uniformity in $P_{\rm eff}$ across density, prevents gas from changing phase and fragmenting significantly.

(vii) We make predictions for the gamma-ray emission from the hadronic pion decay. We find that the time averaged emission in most of our SN placement models with the fiducial diffusivity of 10^{28} cm^2 s^{-1} matches the MW emission (see Fig. 15). Variability in the gamma-ray emission is driven by the proximity of SN explosions to dense gas, which we demonstrate with our models that explicitly vary this.

(viii) We show that the presence of CRs enhances the volume density of neutral atomic hydrogen at heights several 100 pc above the mid-plane in all SN placement models, as shown in Fig. 17. The scale heights of H I volume density is consistent with the scale heights of 21 cm emission when considering most of the mid-plane H I simulated here will be optically thick.

(ix) We examined the volumetric partition of gas of different temperatures and found that in all the SN placement models we have tested, it is difficult to sustain a volume fraction of hot gas of 50 per cent. This may be due to physical effects not considered here such as radiation transport or variable CR diffusion or streaming.

The stratified box model tested here is a flexible platform on which to conduct parameter space studies and test physical ideas in a controlled fashion as we have done here. We have clearly

demonstrated that CR transport and pressure are important effects that should be included in studies of the ISM going forward. We plan to extend this work to also explore how CRs combine with radiation transport in the future, along with varying CR transport (including streaming and diffusion), and effects from runaway stars.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

Ackermann M. et al., 2012, ApJ, 755, 164
Andersson E. P., Agertz O., Renaud F., 2020, MNRAS, 494, 3328
Bell A. R., 2004, MNRAS, 353, 550
Booth C. M., Agertz O., Kravtsov A. V., Gnedin N. Y., 2013, ApJ, 777, L16
Boulares A., Cox D. P., 1990, ApJ, 365, 544
Buck T., Pfrommer C., Pakmor R., Grand R. J. J., Springel V., 2020, MNRAS, 497, 1712
Butsky I. S., Quinn T. R., 2018, ApJ, 868, 108
Caprioli D., Spitkovsky A., 2014, ApJ, 783, 91
Caprioli D., Haggerty C. C., Blasi P., 2020, ApJ, 905, 2
Clark P. C., Glover S. C. O., Klessen R. S., 2012, MNRAS, 420, 745
Commerçon B., Marcowith A., Dubois Y., 2019, A&A, 622, A143
Cox D. P., 2005, ARA&A, 43, 337
Creasey P., Theuns T., Bower R. G., 2013, MNRAS, 429, 1922
Dashyan G., Dubois Y., 2020, A&A, 638, A123
Draine B. T., 1978, ApJS, 36, 595
Draine B. T., 2011, Physics of the Interstellar and Intergalactic Medium. Princeton Univ. Press, Princeton
Enßlin T. A., Pfrommer C., Springel V., Jubelgas M., 2007, A&A, 473, 41
Evoli C., Gaggero D., Grasso D., Maccione L., 2008, J. Cosmol. Astropart. Phys., 2008, 018
Evoli C., Morlino G., Blasi P., Aloisio R., 2020, Phys. Rev. D, 101, 023013
Farber R., Ruszkowski M., Yang H. Y. K., Zweibel E. G., 2018, ApJ, 856, 112
Farcy M., Rosdahl J., Dubois Y., Blaizot J., Martin-Alvarez S., 2022, MNRAS, 513, 5000
Flynn C., Holmberg J., Portinari L., Fuchs B., Jahreiß H., 2006, MNRAS, 372, 1149
Fox A. J., Richter P., Ashley T., Heckman T. M., Lehner N., Werk J. K., Bordoloi R., Peeples M. S., 2019, ApJ, 884, 53
Girichidis P. et al., 2016, ApJ, 816, L19
Girichidis P., Naab T., Hanasz M., Walch S., 2018, MNRAS, 479, 3042
Girichidis P., Pfrommer C., Hanasz M., Naab T., 2020, MNRAS, 491, 993
Girichidis P., Pfrommer C., Pakmor R., Springel V., 2022, MNRAS, 510, 3917
Glover S. C. O., Clark P. C., 2012, MNRAS, 421, 116
Glover S. C. O., Mac Low M.-M., 2007a, ApJS, 169, 239
Glover S. C. O., Mac Low M.-M., 2007b, ApJ, 659, 1317
Gnat O., Ferland G. J., 2012, ApJS, 199, 20
Guo F., Oh S. P., 2008, MNRAS, 384, 251
Habing H. J., 1968, Bull. Astron. Inst. Netherlands, 19, 421

Haffner L. M., Reynolds R. J., Tufte S. L., Madsen G. J., Jaehnig K. P., Percival J. W., 2003, ApJS, 149, 405
Haggerty C. C., Caprioli D., 2020, ApJ, 905, 1
Hill A. S., Joung M. R., Mac Low M.-M., Benjamin R. A., Haffner L. M., Klingenberg C., Waagan K., 2012, ApJ, 750, 104
Hopkins P. F., Butsky I. S., Panopoulou G. V., Ji S., Quataert E., Faucher-Giguere C.-A., Keres D., 2021a, MNRAS, 516, 3470
Hopkins P. F., Chan T. K., Ji S., Hummels C. B., Kereš D., Quataert E., Faucher-Giguère C.-A., 2021b, MNRAS, 501, 3640
Hopkins P. F., Squire J., Butsky I. S., Ji S., 2022, MNRAS, 517, 5413
Iffrig O., Hennebelle P., 2015, A&A, 576, A95
Jacob S., Pakmor R., Simpson C. M., Springel V., Pfrommer C., 2018, MNRAS, 475, 570
Jiang Y.-F., Oh S. P., 2018, ApJ, 854, 5
Joshi P. R., Walch S., Seifried D., Glover S. C. O., Clarke S. D., Weis M., 2019, MNRAS, 484, 1735
Jubelgas M., Springel V., Enßlin T., Pfrommer C., 2008, A&A, 481, 33
Kalberla P. M. W., Dedes L., 2008, A&A, 487, 951
Kannan R., Marinacci F., Simpson C. M., Glover S. C. O., Hernquist L., 2020, MNRAS, 491, 2088
Kennicutt R. C., Jr, 1998, ApJ, 498, 541
Kissmann R., Werner M., Reimer O., Strong A. W., 2015, Astropart. Phys., 70, 39
Krumholz M. R., Burkhart B., 2016, MNRAS, 458, 1671
Kulsrud R., Pearce W. P., 1969, ApJ, 156, 445
Lacki B. C., Thompson T. A., Quataert E., 2010, ApJ, 717, 1
Lloyd S., 1982, IEEE Trans. Inform. Theory, 28, 129
Long K. S., Blair W. P., Winkler P. F., Della Bruna L., Adamo A., McLeod A. F., Amram P., 2022, ApJ, 929, 144
Machacek M. E., Bryan G. L., Abel T., 2001, ApJ, 548, 509
Marinacci F., Vogelsberger M., 2016, MNRAS, 456, L69
McKee C. F., Ostriker J. P., 1977, ApJ, 218, 148
Morlino G., Caprioli D., 2012, A&A, 538, A81
Moskalenko I. V., Strong A. W., 1998, ApJ, 493, 694
Naab T., Ostriker J. P., 2017, ARA&A, 55, 59
Nelson R. P., Langer W. D., 1997, ApJ, 482, 796
Ohlin L., Renaud F., Agertz O., 2019, MNRAS, 485, 3887
Pais M., Pfrommer C., Ehlert K., Pakmor R., 2018, MNRAS, 478, 5278
Pakmor R., Springel V., 2013, MNRAS, 432, 176
Pakmor R., Springel V., Bauer A., Mocz P., Munoz D. J., Ohlmann S. T., Schaal K., Zhu C., 2016a, MNRAS, 455, 1134
Pakmor R., Pfrommer C., Simpson C. M., Kannan R., Springel V., 2016b, MNRAS, 462, 2603
Pakmor R. et al., 2017, MNRAS, 469, 3185
Pfrommer C., Enßlin T. A., 2004, A&A, 413, 17
Pfrommer C., Pakmor R., Schaal K., Simpson C. M., Springel V., 2017a, MNRAS, 465, 4500
Pfrommer C., Pakmor R., Simpson C. M., Springel V., 2017b, ApJ, 847, L13
Pfrommer C., Werhahn M., Pakmor R., Girichidis P., Simpson C. M., 2021, MNRAS, 515, 4229
Quataert E., Thompson T. A., Jiang Y.-F., 2022, MNRAS, 510, 1184
Rathjen T.-E. et al., 2021, MNRAS, 504, 1039
Rojas-Bravo C., Araya M., 2016, MNRAS, 463, 1068
Ruszkowski M., Yang H. Y. K., Zweibel E., 2017, ApJ, 834, 208
Salem M., Bryan G. L., 2014, MNRAS, 437, 3312
Salem M., Bryan G. L., Hummels C., 2014, ApJ, 797, L18
Schroer B., Pezzi O., Caprioli D., Haggerty C. C., Blasi P., 2022, MNRAS, 512, 233
Seifried D. et al., 2017, MNRAS, 472, 4797
Sembach K. R., Howk J. C., Ryans R. S. I., Keenan F. P., 2000, ApJ, 528, 310
Semenov V. A., Kravtsov A. V., Caprioli D., 2021, ApJ, 910, 126
Shalaby M., Thomas T., Pfrommer C., 2021, ApJ, 908, 206
Shukurov A., Sokoloff D., Subramanian K., Brandenburg A., 2006, A&A, 448, L33
Simpson C. M., Bryan G. L., Hummels C., Ostriker J. P., 2015, ApJ, 809, 69
Simpson C. M., Pakmor R., Marinacci F., Pfrommer C., Springel V., Glover S. C. O., Clark P. C., Smith R. J., 2016, ApJ, 827, L29

Smith R. J., Glover S. C. O., Clark P. C., Klessen R. S., Springel V., 2014, MNRAS, 441, 1628

Springel V., 2010, MNRAS, 401, 791

Springel V., Pakmor R., Zier O., Reinecke M., 2021, MNRAS, 506, 2871

Thomas T., Pfrommer C., 2019, MNRAS, 485, 2977

Thomas T., Pfrommer C., 2022, MNRAS, 509, 4803

Thomas T., Pfrommer C., Enßlin T., 2020, ApJ, 890, L18

Thomas T., Pfrommer C., Pakmor R., 2021, MNRAS, 503, 2242

Thomas T., Pfrommer C., Pakmor R., 2022, preprint (arXiv:2203.12029)

Thompson T. A., Quataert E., Waxman E., Murray N., Martin C. L., 2006, ApJ, 645, 186

Truelove J. K., Klein R. I., McKee C. F., Holliman, John H. I., Howell L. H., Greenough J. A., 1997, ApJ, 489, L179

Tumlinson J., Peeples M. S., Werk J. K., 2017, ARA&A, 55, 389

Uhlig M., Pfrommer C., Sharma M., Nath B. B., Enßlin T. A., Springel V., 2012, MNRAS, 423, 2374

Walch S., Naab T., 2015, MNRAS, 451, 2757

Walch S. et al., 2015, MNRAS, 454, 238

Wechsler R. H., Tinker J. L., 2018, ARA&A, 56, 435

Werhahn M., Pfrommer C., Girichidis P., Puchwein E., Pakmor R., 2021a, MNRAS, 505, 3273

Werhahn M., Pfrommer C., Girichidis P., Winner G., 2021b, MNRAS, 505, 3295

Werhahn M., Pfrommer C., Girichidis P., 2021c, MNRAS, 508, 4072

Werk J. K. et al., 2019, ApJ, 887, 89

Wiener J., Pfrommer C., Oh S. P., 2017, MNRAS, 467, 906

Zweibel E. G., 2017, Phys. Plasmas, 24, 055402

APPENDIX A: PRESSURE LIMITER

Here we describe how the Riemann solver pressure limiter impacts the behaviour of the model. This model introduces a variable parameter ρ_{max} which is the maximum density of a Jeans stable cloud (see equation 3). In gas above this density, the gravitational self-collapse of gas is prevented by a limit placed on the pressure used in the Riemann solver, thereby limiting the fluxes between cells. Without this limiter, gas will collapse to arbitrarily high density because of gas self-gravity and radiative cooling.

The collapse of gas to a small volume can be prevented by halting cell refinement at a minimum cell volume, a common approach in adaptive mesh refinement (AMR) codes and one employed in Simpson et al. (2016). However, we found that the extreme gradients resulting from overmassive cells produced spurious amplification of the simulated magnetic field, and made this an unworkable solution for our problem.

Other similar simulations have side-stepped this issue by not including gas self-gravity (e.g. Hill et al. 2012), and indeed this limits the collapse of clouds, but may impact the fragmentation of the medium on certain scales (e.g. Walch et al. 2015). Another approach is to provide a sink for collapsing gas in the form of star particles that will shift gas mass to a collision-less component that will be subject to gravitational softening. However, with the resolution we employ here, only a few per cent of collapsing gas should go into stars, making this an inappropriate solution to halt collapse. Ultimately, neglecting low-temperature cooling below 10^4 K would also be effective in preventing the collapse of dense gas (e.g. Creasey et al. 2013), but of course, this does not allow for the study of the structure of the ISM on sub-kiloparsec scales as is our goal.

The main goal the simulations presented here is to understand the impact of CRs on the scale of the local solar neighbourhood, a region of roughly a kpc in scale, and not to study the evolution of molecular clouds, for which we do not include necessary physical effects or appropriate resolution.
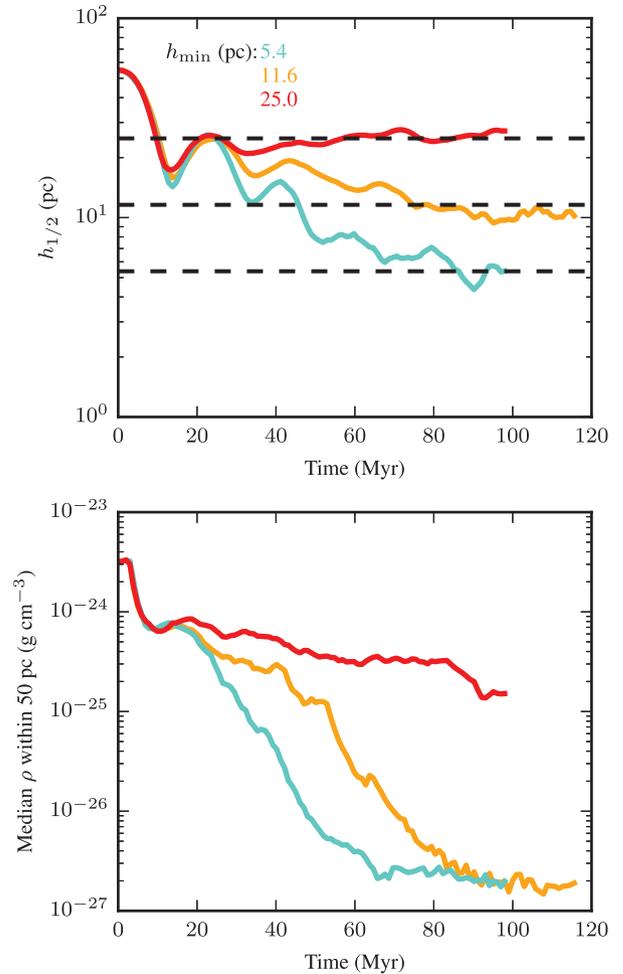
Figure A1. Evolution of three boxes simulated with only hydrodynamics (no CRs and no magnetic fields) where 50 per cent of SN explosions occur in random locations and 50 per cent take place in high-density gas (model H0.5 in the text). These models all have different values of ρ_{max}: the fiducial value used in the main text of $\rho_{max} \sim 2 \times 10^{-22}$ (gold), a ρ_{max} ten times the fiducial value (cyan) and one tenth of the fiducial ρ_{max} (red). The top panel shows the height containing half the original mass of the box and the bottom panel shows the volume-weighted median gas density within 50 pc of the mid-plane. The top panel includes dashed lines showing the radii of a Jeans stable cloud with average densities at these different values of ρ_{max}.
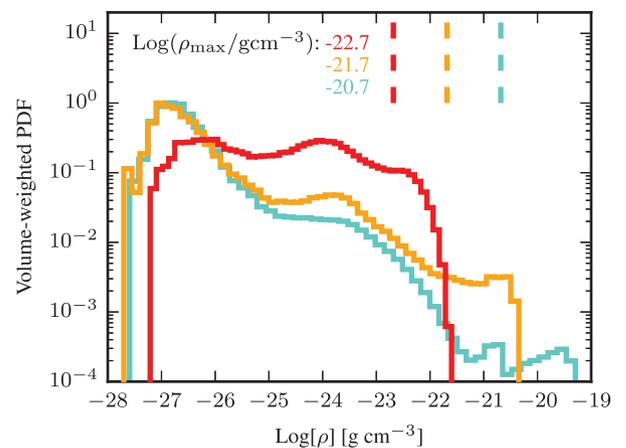
Figure A2. The volume-weighted distribution of gas densities after 100 Myr of evolution for the three models shown in Fig. A1. The value of ρ_{max} for each model is indicated.
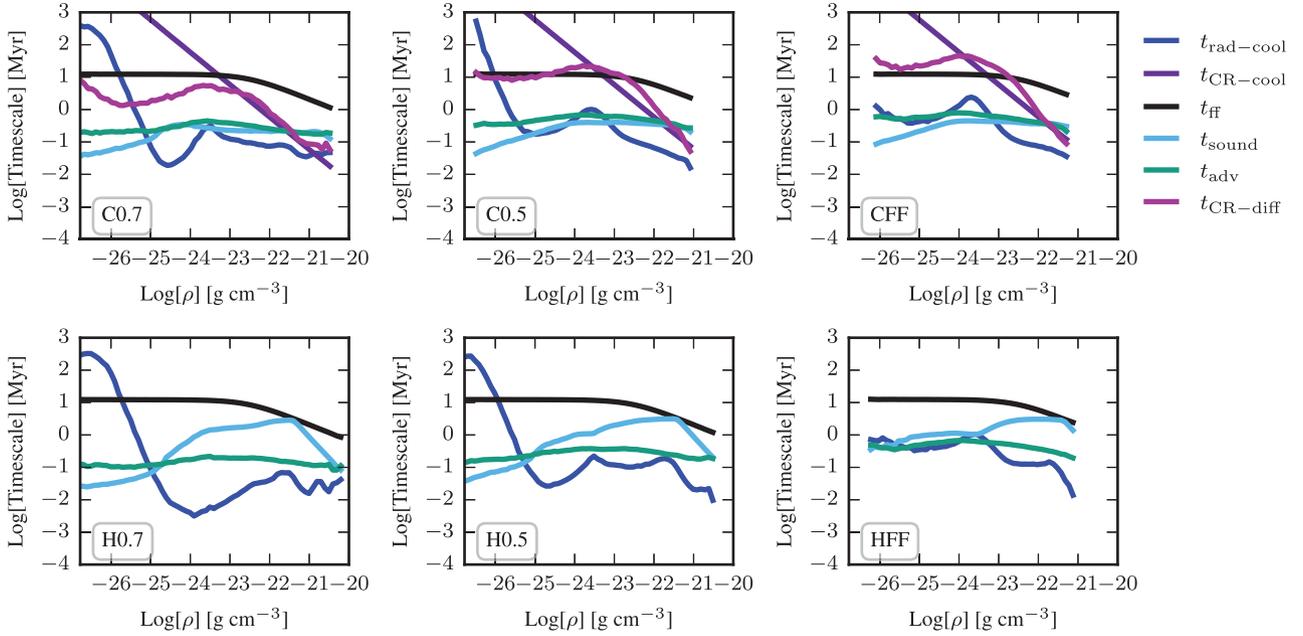
Figure A3. Time-scales of the ISM within 50 pc of the mid-plane after 100 Myr of evolution in models with varying SN explosion placement (left to right) and with and without CRs (top and bottom). The median time-scale is plotted as a function of gas density. The time-scales presented are the effective cooling time that accounts for CRs and their hadronic cooling ($t_{\text{eff-cool}}$), the gravitational free fall time (t_{ff}), the sound crossing time (t_{sound}), the turbulent crossing time (t_{turb}), the CR diffusion time ($t_{\text{CR-diff}}$), and the CR cooling time ($t_{\text{CR-cool}}$). The models C0.7, H0.7, and H0.5 have undergone a thermal runaway and their mid-plane volume is dominated by hot, low density gas. The models C0.5, CFF, and HFF remain in the peak driving solution and their mid-plane volume is dominated by warm gas. The transition between the warm and cold phase of the ISM occurs around 10^{23} g cm^{-3} and the time-scales around this density are the most relevant to the problem.

Fig. A1 shows the behaviour of test runs with a higher and lower values of ρ_{max} at fixed mass resolution ($m_t = 10 \, \text{M}_\odot$). In effect, what varies in these models is the radius of a Jeans stable cloud in the simulation, where runs with low values of ρ_{max} are only able to capture larger clouds. Fig. A1 shows that indeed, gas in these runs collapses to a layer with a height equivalent to the size of a Jeans stable cloud of this mass.

This model is intended to halt the collapse of dense clouds in the absence of a full model for star formation and pre-SN feedback. However, if applied in extremity, this pressure limiter can act as a form of subgrid feedback that impacts gas at densities below ρ_{max}. Fig. A2 shows this effect. The model with the lowest ρ_{max} tested has more gas at densities of 10^{-24} g cm^{-3} because the low value of ρ_{max} prevents gas at these densities from collapsing to higher densities. Our fiducial value of ρ_{max} allows for this collapse and results in a gas density distribution resembling that of a model with a higher ρ_{max} but lacking the the tail of densities above 10^{-20} g cm^{-3}.

These models' large scale behaviour largely depends on whether gas can transition from the WIM around a temperature of 10^4 down to temperatures around 100 K of the cold neutral medium. Fig. 1 shows this transition happens around a density of 10^{-23} g cm^{-3}. A choice of ρ_{max} too close to this value will impede this transition and affect the development of the WNM and CNM. This in turn will impact the porosity of the medium and the amount of gas available to be driven from the mid-plane in outflows.

We have presented tests with a feedback model where 50 per cent of SN explosions occur in random locations and 50 per cent in dense gas. With our fiducial model parameters, this model undergoes a thermal runaway. This model without CRs is the model closest to tipping over to the peak driving solution. It behaves similarly to the model with a value of ρ_{max} 10 times larger. They reach a similar

average log-density and have similar gas density distributions up to densities of ρ_{max}. With a value of ρ_{max} 0.1 times smaller, the medium does not undergo the same type of collapse. However, the distribution of gas is different from the peak driving models seen in Fig. 11. Gas is more extended in height by a factor of several and it's distribution in density is broader, but more uniform up to ρ_{max}.

The choice of ρ_{max}, therefore does have the ability to significantly alter the model's behaviour, but here we have demonstrated our fiducial value of ρ_{max} for our chosen mass resolution allows this transitional case to still undergo a thermal runaway. Of course, in models with a large number of SN explosions occurring in dense gas, the injected SN energy make the pressure limiter irrelevant because the injected energy is effective in countering cooling and prevents gas from reaching ρ_{max}. Models that have a high f_{rand} are most sensitive to the details of the pressure limiter.

APPENDIX B: TIME-SCALES

Useful quantities to consider in the evolution of the ISM are the time-scales of various gas processes shown in Fig. A3. The time-scales we consider are the gravitational free fall time (t_{ff}, see equation 13), the radiative cooling time of the thermal gas ($t_{\text{rad-cool}}$), the sound crossing time (t_{sound}), the advective crossing time (t_{adv}), the CR cooling time ($t_{\text{CR-cool}}$), the CR diffusion time ($t_{\text{CR-diff}}$), and the SNe heating time ($t_{\text{SN-heat}}$).

The radiative cooling time is an output of the chemistry and cooling network used here (Glover & Mac Low 2007a, b; Smith et al. 2014) that includes atomic and molecular hydrogen cooling along with metal line cooling for a metallicity of $1 \, \text{Z}_\odot$.

The CR cooling time is defined as the time it would take to reduce the CR energy in a cell by an e-folding as a result of hadronic and

Coulomb losses:

$$t_{\mathrm{CR-cool}} = \frac{\varepsilon_{\mathrm{CR}}\rho}{\Lambda_{\mathrm{hadr}} + \Lambda_{\mathrm{Coul}}}, \tag{B1}$$

where Λ_{Coul} and Λ_{hadr} are the volumetric loss rates of CR energy from these processes (assuming a steady-state CR distribution).

The CR diffusion time is related to the CR pressure gradient in the direction of the local magnetic field (**b**) and the diffusivity κ:

$$t_{\mathrm{CR-diff}} = \left(\frac{P_{\mathrm{CR}}}{\mathbf{b}\cdot\nabla P_{\mathrm{CR}}}\right)^2 \kappa^{-1}. \tag{B2}$$

We also consider the sound crossing time and the advective crossing time, which are the ratios of the cell's approximate diameter Δx derived from its volume V_i and the sound speed of the medium in the cell ($c_{i,\mathrm{s}}$) or the bulk velocity of the cell (v_i), respectively:

$$t_{\mathrm{sound}} = \Delta x c_{i,\mathrm{s}}^{-1} \tag{B3}$$

and

$$t_{\mathrm{adv}} = \Delta x v_i^{-1}, \tag{B4}$$

where $\Delta x = (6V_i/\pi)^{1/3}$. The sound speed depends on the pressure within the cell:

$$c_{i,\mathrm{s}} = \left(\frac{\frac{5}{3}P_{i,\mathrm{therm}} + \frac{4}{3}P_{i,\mathrm{CR}} + B_i^2/(8\pi)}{\rho_i}\right)^{1/2}, \tag{B5}$$

where $P_{i,\mathrm{therm}}$ is the thermal pressure and $P_{i,\mathrm{CR}}$ is the CR pressure. Models that do not include magnetic fields or CRs only include the thermal pressure term. In the case of non-zero-magnetic fields (that we evolve with the ideal MHD assumption), we adopt the maximum speed of the fast magnetosonic wave that is driven by magnetic pressure. Models with CRs tend to have a higher total pressures due to the contribution of the CR fluid. This results in higher sound speeds and lower sound crossing times.

Fig. A3 compares these time-scales as a function of density in runs with different SN placement models and with and without CRs. The density where gas transitions from the warm phase to the cold phase is around $10^{23}\,\mathrm{g\,cm^{-3}}$ as shown in Fig. 1.

In all models at the relevant densities, the median free-fall time is typically too long to be relevant, but that does not mean that on small scales gas self-gravity cannot play a role in the collapse of gas. What is perhaps more interesting is the role of gas motions, which can be a source of dynamical heating, but can also cause the collapse of gas through shock compression or shock collisions. The sound crossing time becomes shorter with the addition of CRs due to their impact on the sound speed in the gas (it becomes faster). The cell-crossing (or advection) time remains unchanged, indicating that the typical gas bulk motions remain unchanged with the addition of CRs.

The CR diffusion time becomes shorter in denser gas. In models with a thermal runaway and very dense gas, it becomes shorter than the radiative cooling time. In this situation, the CR diffusion removes CR energy from dense gas faster than the gas has a chance to cool, indicating CRs leave before they can impact the thermal gas evolution.

This paper has been typeset from a TEX/LATEX file prepared by the author.