



Cosmic-Ray Transport, Energy Loss, and Influence in the Multiphase Interstellar Medium

Chad Bustard^{1,2}  and Ellen G. Zweibel^{1,3} 

¹ Department of Physics, University of Wisconsin-Madison, 1150 University Avenue, Madison, WI 53706, USA; bustard@ucsb.edu

² Kavli Institute for Theoretical Physics, University of California - Santa Barbara, Kohn Hall, Santa Barbara, CA 93107, USA

³ Department of Astronomy, University of Wisconsin - Madison, 475 North Charter Street, Madison, WI 53706, USA

Received 2020 December 11; revised 2021 March 30; accepted 2021 April 7; published 2021 June 1

Abstract

The bulk propagation speed of GeV-energy cosmic rays is limited by frequent scattering off hydromagnetic waves. Most galaxy evolution simulations that account for this confinement assume the gas is fully ionized and cosmic rays are well coupled to Alfvén waves; however, multiphase density inhomogeneities, frequently underresolved in galaxy evolution simulations, induce cosmic-ray collisions and ionization-dependent transport driven by cosmic-ray decoupling and elevated streaming speeds in partially neutral gas. How do cosmic rays navigate and influence such a medium, and can we constrain this transport with observations? In this paper, we simulate cosmic-ray fronts impinging upon idealized, partially neutral clouds and lognormally distributed clumps, with and without ionization-dependent transport. With these high-resolution simulations, we identify cloud interfaces as crucial regions where cosmic-ray fronts can develop a stairstep pressure gradient sufficient to collisionlessly generate waves, overcome ion–neutral damping, and exert a force on the cloud. We find that the acceleration of cold clouds is hindered by only a factor of a few when ionization-dependent transport is included, with additional dependencies on magnetic field strength and cloud dimensionality. We also probe how cosmic rays sample the background gas and quantify collisional losses. Hadronic gamma-ray emission maps are qualitatively different when ionization-dependent transport is included, but the overall luminosity varies by only a small factor, as the short cosmic-ray residence times in cold clouds are offset by the higher densities that cosmic rays sample.

Unified Astronomy Thesaurus concepts: Cosmic rays (329); Magnetic fields (994); Interstellar medium (847); Gamma-rays (637); Galaxy evolution (594)

Supporting material: animations

1. Introduction

Highly energetic cosmic rays, despite representing only a microscopic fraction of particles by number density ($n_{\text{CR}}/n_{\text{gas}} \approx 10^{-9}$), are roughly in energy equipartition with thermal and magnetic energy in the Milky Way interstellar medium (ISM; Boulares & Cox 1990). Because hadronic cosmic rays do not suffer from radiative losses (though hadronic collisions, Coulomb collisions, and “collisionless” scattering off magnetic perturbations can be large energy sinks), and they do not lose as much energy to adiabatic expansion compared to thermal gas ($P \propto \rho^{4/3}$ for a relativistic gas instead of $P \propto \rho^{5/3}$ for a nonrelativistic gas), they fundamentally alter ISM dynamics and are believed to be an important component of feedback, which regulates star formation in galaxies.

For GeV cosmic rays, which make up the peak of the cosmic-ray energy spectrum and therefore contain most of the energy and momentum of the cosmic-ray population, the dominant transport mode is self-confinement via the streaming instability (Wentzel 1968; Kulsrud & Pearce 1969; Wentzel 1969). In this picture, cosmic rays with a bulk drift speed greater than the Alfvén speed can excite Alfvén waves through gyroresonance and pitch angle scatter off the waves until the cosmic rays attain isotropy in the wave frame. Scattering transfers energy from the cosmic-ray population to the waves, and subsequent wave damping heats the background gas. When the scattering mean free path is short⁴, cosmic rays are well coupled to waves, and the bulk population advects at the

Alfvén speed down the cosmic-ray pressure gradient directed along the local magnetic field. Extrinsic turbulence can alternatively confine cosmic rays, in which case cosmic rays scatter off the turbulent cascade. The resulting transport is field-aligned diffusion, but unlike the self-confinement model, there is no net transfer of energy to the thermal gas (Zweibel 2017).

To date, most galaxy evolution simulations that include a relativistic cosmic-ray fluid assume some combination of diffusive and streaming transport in addition to cosmic-ray advection with the nonrelativistic thermal gas. Given the canonical picture that 10% of supernova energy is converted to cosmic-ray energy through first-order Fermi acceleration, simulations with diffusive or streaming transport included generally find colder, smoother, more extended outflows than thermally driven winds, and in many cases, the nonthermal support is necessary to launch an outflow (e.g., Ipavich 1975; Breitschwerdt et al. 1991; Everett et al. 2008; Uhlig et al. 2012; Hanasz et al. 2013; Salem & Bryan 2014; Ruszkowski et al. 2017; Mao & Ostriker 2018; Buck et al. 2020; Dashyan & Dubois 2020; Bustard et al. 2020; Hopkins et al. 2021a). Theories and observations alike also suggest the existence of a long-lived cosmic-ray reservoir in the circumgalactic medium (CGM), where they may further influence the gas cycle in and out of galaxies (Salem & Bryan 2014; Salem et al. 2016; Girichidis et al. 2018; Butsky & Quinn 2018; Blasi & Amato 2019; Heintz et al. 2020; Ji et al. 2020; Kempfski & Quataert 2020), but the extent to which cosmic rays navigate through the disk–halo interface, carry mass-loaded winds, and influence the CGM are transport dependent. Given the myriad possibilities for cosmic rays to narrow gaps between

⁴ In fully ionized Milky Way ISM conditions, this is a good approximation, because the mean free path is ≈ 0.1 pc.

observations and simulations of galaxy evolution, more work is needed to extend and apply first-principles fluid models of cosmic rays to galaxy evolution simulations and, likewise, to constrain these models with mock observations.

An important but mostly overlooked complication, worthy of more detailed study, is that the waves that scatter cosmic rays are not always guaranteed to be present. In particular, density inhomogeneities, frequently underresolved in simulations of galaxy evolution, can induce alternating regions where cosmic rays are coupled or decoupled from waves. Let us consider a front of self-confined cosmic rays impinging upon a cloud. As the cosmic rays stream down their pressure gradient and encounter a drop in Alfvén speed in the cloud, their pressure buildup at the cloud interface exerts a force on the cloud (Wiener et al. 2017a, 2019). Upstream of this “bottleneck”, however, the steady-state solution is a flat cosmic-ray pressure profile unable to excite hydromagnetic waves and unable to exert a force along the direction of the magnetic field (Skilling 1971). Cosmic rays similarly become decoupled when the confining waves are heavily damped, forcing steady-state cosmic-ray streaming speeds to be much higher than the gas Alfvén speed in order for the wave excitation rate ($\propto v_{st}$) to be comparable to the damping rate. In partially neutral clouds, where damping due to ion–neutral friction can be strong, this effect is believed to be especially significant (Kulsrud & Pearce 1969). Additionally, even when cosmic rays are coupled in partially neutral gas, the wave velocity that they move at is the ion Alfvén speed, $v_A^{\text{ion}} = v_A / \sqrt{f_{\text{ion}}}$, which may be much larger than the gas Alfvén speed in media with low ion fraction, f_{ion} . For the remainder of this paper, we will use the term “ionization-dependent” transport to refer to these collective modifications to the standard streaming transport picture.

As it has been shown that cosmic-ray bottlenecks can theoretically accelerate warm, fully ionized clouds in galaxy halos up to 100 s of km s^{−1}, it is of great interest to understand the conditions in the ISM rather than in the CGM, where cosmic rays are similarly well coupled and capable of doing work through their pressure gradients. This problem is not trivial and may have important consequences for feedback.

Farber et al. (2018) ran magnetohydrodynamic (MHD) ISM “patch” simulations of cosmic-ray-driven winds with a locally adaptive diffusion coefficient boosted in cold, predominantly neutral gas to mock up the effects of ion–neutral damping. This led to different gas properties, a broader spatial distribution of cosmic rays, and higher wind speeds as cosmic rays preferentially pushed on the hot gas and escaped more easily into the halo. These pioneering “two- κ ” models, in which the diffusion coefficient shifted between a low and high value depending on gas temperature, underscored the possible implications of decoupling, but the transition from fully ionized to partially neutral transport is highly nonlinear with further dependencies on magnetic field strength and cosmic-ray pressure gradient (Everett & Zweibel 2011). An idealized study focused on teasing out some of the complicated interplay between bottlenecks, variable transport, and collisional losses in the multiphase ISM is still needed and will be presented in this paper.

This study is further motivated by recent full-galaxy simulations from the FIRE collaboration (Chan et al. 2019; Hopkins et al. 2021b) that attempt to constrain cosmic-ray transport with gamma-ray emission and grammage estimates. Intriguingly, the best-fit simulations necessitate very fast

cosmic-ray transport speeds, especially in diffuse gas, to lower gamma-ray emission to observed values, while fast transport just in cold, partially neutral clouds is insufficient. Such a high gamma-ray overproduction (a factor of 10 or greater) in simulated dwarf galaxies, also found in LMC-specific simulations of Bustard et al. (2020), demands a follow-up on smaller scales. In this paper, we do not include gravity or feedback, but this step back in complexity is necessary in order to focus on the typically unresolved and already complex cosmic-ray–cloud interactions without confounding results with changes in star formation rates, etc.

Utilizing the Jiang & Oh (2018) cosmic-ray framework implemented in the Athena++ MHD code, we run high-resolution (parsec to subparsec) 1D, 2D, and 3D simulations of cosmic-ray fronts running through various multiphase “obstacle courses”—either single clouds or a lognormal distribution of clumps—to probe cosmic-ray transport speeds, energy loss, and momentum transfer in idealized multiphase ISM environments, focusing on transport with and without neutral gas effects included.

The paper is outlined as follows. In Section 2, we provide more background on the nonlinear cosmic-ray transport and energy-loss mechanisms we study in this work. We present our assumptions and analytic expectations for cosmic-ray transport in Section 3 and our implementation of these effects in the Athena++ MHD code (Stone et al. 2020) in Section 4. In Section 5, we present and compare 1D and 2D simulations of cosmic-ray fronts impinging upon representative ISM clouds. In Section 6, we present 2D and 3D simulations of cosmic-ray transport through a layer of lognormally distributed multiphase clumps in a mock ISM. For each “obstacle course”, we quantify the cosmic-ray energy loss (both collisional and collisionless), the momentum imparted to the gas, and probability distribution functions of cosmic-ray energy and collisional energy loss (with and without taking into account neutral gas). In Section 7, we discuss the results and limitations of our work, and we conclude in Section 8.

Given the range of applications covered in this paper, the reader who is primarily interested in cosmic-ray energy loss and implications for gamma-ray observations will want to read Section 6. Readers who are interested in cosmic-ray influence on galaxy evolution should also read Section 5.

Simulation visualizations can be found at <https://bustardchad.wixsite.com/mysite/visualizations>.

2. Background

2.1. Cosmic-Ray Coupling and Decoupling

To first order, cosmic rays comprise a second, relativistic fluid that advects with the nonrelativistic thermal gas. We additionally know from the near isotropy of observed cosmic rays and their long confinement times in the galaxy, estimated from spallation measurements, that cosmic rays are constantly being scattered (see recent reviews by, e.g., Zweibel 2017; Amato & Blasi 2018; Becker Tjus & Merten 2020). Their resulting transport is well described by a random walk along magnetic field lines with an energy-dependent diffusion coefficient. For cosmic rays at energies above a few hundred GeV, scattering off waves generated by an external turbulent cascade likely dominates the transport. Cosmic rays are coupled to the waves through resonant interactions, and transport is akin to magnetic-field-aligned diffusion with a

coefficient depending on the amplitude of scattering waves. Unless the left- and right-propagating Alfvén waves in the turbulent cascade have unequal intensities, there is no net transfer of energy between cosmic rays and the thermal gas.

For GeV-energy cosmic rays, however, the confining waves can be generated by the cosmic rays themselves. Cosmic rays with even a tiny amount of drift anisotropy can resonate with magnetic perturbations and become self-confined: they exchange energy with Alfvén waves through gyroresonance, pitch angle scatter off those waves, and become locked to the wave frame, therefore advecting at the local ion Alfvén speed $v_A^{\text{ion}} = B/\sqrt{4\pi\rho_{\text{ion}}}$ (in addition to advecting with the background gas flow). In a steady state, the conversion of cosmic-ray energy to wave energy is balanced by wave damping, which heats the background gas at a rate $\propto v_A^{\text{ion}} \cdot \nabla P_{\text{CR}}$. We refer to this energy transfer as collisionless because it is mediated by magnetic fields, not direct collisions between particles.

In a multiphase ISM punctuated by density irregularities, however, cosmic-ray coupling can quickly break down. For instance, consider a decrease in Alfvén speed along a magnetic field line. One can show that, in steady state, streaming cosmic rays must conform to a constant cosmic-ray pressure profile for an extended region upstream (Skilling 1971). Under such conditions, no waves are excited⁵, and cosmic rays no longer transfer energy or momentum to the thermal gas. We will refer to these regions as cosmic-ray “bottlenecks” (Wiener et al. 2017a; Zweibel 2017; Wiener et al. 2019), a subset of a larger class of decoupled regions that Skilling (1971) referred to as “free zones.”

This bottleneck effect has been studied with 1D and 2D numerical simulations under simplified conditions, namely a cosmic-ray front impinging on either a 1D slab or 2D cylindrical cloud. While these simulations promisingly show that cosmic-ray bottlenecks may actually accelerate these cold clouds (Wiener et al. 2017a, 2019; Brüggén & Scannapieco 2020), thereby helping entrain them into hot outflows, they most importantly expose the complex relationship between cosmic rays and the multiphase ISM. Although cosmic-ray transport is determined by local conditions on scales of a cosmic-ray gyroradius, the global environment determines these local conditions. In the bottleneck scenario, the presence of a cold cloud creates a traffic jam, decoupling cosmic rays from thermal gas, extending for macroscopic distances upstream of the cloud. One can then imagine a checkerboard of these clouds, each triggering cosmic-ray bottlenecks. How a steady source of cosmic rays navigates this series of traffic jams depends on the global cloud and magnetic field structure.

Another example of “free zones” can occur in colder, partially neutral gas, where ion–neutral collisions decouple ions and neutrals at scales above the cosmic-ray gyroradius and thereby damp the Alfvén waves that play the crucial scattering role (Kulsrud & Pearce 1969). In a steady state, the resonant streaming instability growth rate, which is proportional to the cosmic-ray drift velocity, must balance this increased damping

rate, effectively increasing the streaming velocity (sometimes by orders of magnitude). If ion–neutral damping is strong enough, the cosmic-ray mean free path may in fact be larger than the partially neutral gas region, leading cosmic rays to free-stream. Even if the cosmic rays can drive waves at a rate fast enough to couple them to the gas, cosmic rays only scatter off Alfvén waves that propagate in the ions, meaning the streaming velocity is rigorously the ion Alfvén speed, which can be orders of magnitude larger than the gas Alfvén speed when $\rho_{\text{ion}} \ll \rho$.⁶

2.2. Collisional Losses

In addition to instigating complex, nonlinear transport, cold clouds are also targets for cosmic-ray hadronic and Coulomb collisions, the former leading to a catastrophic decay of pions into gamma-rays. These collisional losses, depending on the environment, can represent large energy sinks that suppress cosmic-ray influence in galaxies (Hopkins et al. 2020c; Crocker et al. 2021). This is most true for starburst galaxies, which have high volume-filling factors of dense, neutral gas and by all inferences from their gamma-ray emission appear to be good proton calorimeters (Lacki et al. 2011; Tang et al. 2014; Yoast-Hull et al. 2015, 2016; Wang & Fields 2018; Krumholz et al. 2020), meaning that all cosmic rays produced from supernovae are consumed by collisions. Observations of L_* and dwarf galaxies, however, paint a different picture where the vast majority of protons escape hadronic collisions presumably due to a combination of diffusion and advection in supernova-driven winds (Lacki et al. 2011; Ackermann et al. 2016; Fu et al. 2017; Lopez et al. 2018). While we only have limited gamma-ray observations of ≈ 10 external star-forming galaxies (e.g., Ajello et al. 2020), and while it is difficult to exactly separate the diffuse hadronic gamma-ray emission from point-source emission, the resulting scaling relations between gamma-ray luminosity and star formation rate (SFR) are already in severe disagreement with the current iteration of MHD + cosmic-ray dwarf and L_* galaxy simulations—that is, simulations that include cosmic-ray diffusive or streaming transport but assume that transport is not affected by ionization.

Chan et al. (2019) and Hopkins et al. (2021b) vary cosmic-ray transport models in their state-of-the-art full-galaxy simulations and find that gamma-rays will be overproduced compared to Local Group dwarf galaxy observations unless the cosmic-ray diffusion coefficient is large ($> 10^{29} \text{ cm}^2 \text{ s}^{-1}$) compared to the canonical value of a few $\times 10^{28} \text{ cm}^2 \text{ s}^{-1}$ inferred from cosmic-ray propagation models. Testing a variety of proposed first-principles transport models, Hopkins et al. (2021b) constrain the issue further: propagation appears to be rate-limited by the warm ionized medium (WIM) and inner CGM, where the current paradigm of self-confinement models predicts relatively slow transport speeds. Fast transport in partially neutral gas is not sufficient to decrease gamma-ray emission, which is instead dominated by small patches of the WIM where the authors believe a runaway increase in self-confinement could be occurring. That is, large cosmic-ray pressure gradients could generate confining waves that further

⁵ Cosmic-ray pressure anisotropy may provide an alternative confinement mechanism by exciting magnetic waves in the upstream region, where the cosmic-ray streaming instability outlined here does not act. This could lock cosmic rays back to the thermal gas (instead of allowing them to decouple), but appears not to transfer much momentum or energy to the gas (Zweibel 2020). For this work, we will only consider cosmic-ray confinement through drift anisotropy and save an analysis of pressure anisotropy for future work.

⁶ For the remainder of this paper, we will distinguish between these two important effects, reserving the frequently used term “fast transport” specifically for elevated diffusive transport arising from cosmic-ray decoupling. We will refer to the combination of decoupling and elevated ion Alfvén speeds when ion fraction $f_{\text{ion}} < 1$ as “ionization-dependent transport”.

trap cosmic rays, increase the cosmic-ray pressure gradient, and so on.

Clearly, the interplay between cosmic-ray transport, wave damping, and collisional losses in the multiphase ISM is very rich and has implications for gamma-ray constraints as well as general ISM and galactic wind dynamics. The goal of this paper is to shine a light on this interplay with idealized but high-resolution simulations.

3. Analytic Expectations

Within the self-confinement picture of cosmic-ray transport, the streaming velocity is

$$\mathbf{v}_{\text{st}} = \frac{\mathbf{B} \cdot (\nabla \cdot \mathbf{P}_{\text{CR}})}{|\mathbf{B} \cdot (\nabla \cdot \mathbf{P}_{\text{CR}})|} v_A^{\text{ion}}, \quad (1)$$

where $v_A^{\text{ion}} = B/\sqrt{4\pi\rho_{\text{ion}}} = B/\sqrt{4\pi\rho f_{\text{ion}}}$, where f_{ion} is the ion fraction by number, and $\rho = mn$ is the total gas density. Here, and in the following shortened derivation (see also Jiang & Oh 2018; Hopkins et al. 2021b), we have been careful to write the Alfvén speed as the ion Alfvén speed, v_A^{ion} , to reflect that cosmic rays resonate with Alfvén waves that propagate in the ions. Note that although Equation (1) is written in terms of the cosmic-ray pressure tensor $\mathbf{P}_{\text{CR}}$, we will assume isotropic pressure and replace the divergence with a gradient. As seen from Equation (1), self-confinement only occurs in the presence of a cosmic-ray pressure gradient directed along the magnetic field. Equation (1) holds, as written, when there is no wave damping, but more generally, the steady-state velocity is obtained by equating the growth rate of cosmic-ray-excited Alfvén waves Γ_{CR} with the wave damping rate Γ :

$$\Gamma_{\text{CR}}(\gamma) \approx \frac{\pi}{4} \frac{\alpha - 3}{\alpha - 2} \Omega_0 \frac{n_{\text{CR}}(>\gamma)}{n_i} \left(\frac{v_D}{v_A^{\text{ion}}} - 1 \right) = \Gamma, \quad (2)$$

where n_i is the ion number density, $n_{\text{CR}}(>\gamma)$ is the number density of cosmic rays with Lorentz factor greater than γ , and we will assume $\alpha = 4$ is the power-law exponent of the cosmic-ray distribution function in momentum space.

Because the growth rate of the streaming instability scales linearly with the bulk drift speed of the cosmic rays, a higher drift speed is needed when wave damping is present. The net drift with respect to the Alfvén wave frame is $v_D - v_A^{\text{ion}}$, and we write this as a diffusive flux (Jiang & Oh 2018):

$$F_{\text{diffuse}} = \kappa_{\parallel} \nabla f \approx (v_D - v_A^{\text{ion}}) f. \quad (3)$$

Then the (purely parallel) diffusivity is

$$\kappa_{\parallel} = \frac{f}{\nabla f} \frac{4}{\pi} \frac{\alpha - 3}{\alpha - 2} \frac{\Gamma n_i v_A^{\text{ion}}}{\Omega_0 n_{\text{CR}}(>\gamma)}. \quad (4)$$

Following Hopkins et al. (2021b), we write $f/\nabla f = l_{\text{CR}}$, a cosmic-ray scale length, and we can formulate this diffusion coefficient as an effective boost to the streaming velocity, $v_{\text{st}}^{\text{eff}} = v_A^{\text{ion}} + \kappa_{\parallel}/(\gamma_{\text{CR}} l_{\text{CR}})$, where $\gamma_{\text{CR}} = 4/3$. Note that in our simulations, we implement the diffusive flux $\kappa_{\parallel}$ and model streaming with speed $v_{\text{st}} = v_A^{\text{ion}}$, not $v_{\text{st}} = v_{\text{st}}^{\text{eff}}$. We then make some substitutions ($e_B = 1/2(v_A^{\text{ion}})^2 n_i m_p$, $e_{\text{CR}} = m_p n_{\text{CR}} c^2$) to rewrite Equation (4) in terms of more typical code quantities to

obtain

$$\kappa_{\parallel} = \frac{4cr_L \Gamma e_B l_{\text{CR}}}{\pi v_A^{\text{ion}} e_{\text{CR}}}, \quad (5)$$

$$v_{\text{st}}^{\text{eff}} = v_A^{\text{ion}} + \frac{3cr_L \Gamma e_B}{\pi v_A^{\text{ion}} e_{\text{CR}}}, \quad (6)$$

where $r_L = c/\Omega$ is the cosmic-ray Larmor radius, and we assume cosmic rays occupy a single energy bin at 1 GeV.

The damping rate Γ can have contributions from multiple processes, including nonlinear Landau damping, which we find to have no bearing on our results (see Section 7) and turbulent damping (Yan & Lazarian 2004; Lazarian 2016; Holguin et al. 2019). In this work, we will focus on ion-neutral damping. The damping rate for a hydrogen-helium plasma is (Kulsrud & Pearce 1969; O’C Drury et al. 1996; De Pontieu et al. 2001; Hopkins et al. 2021b)

$$\Gamma_{\text{in}} = \frac{\nu_{\text{in}}}{2} \approx 10^{-9} f_{\text{neutral}} T_{1000}^{1/2} \rho_{-24} \text{ s}^{-1}, \quad (7)$$

where ν_{in} is the collision frequency between ions and neutrals and $f_{\text{neutral}} = 1 - f_{\text{ion}}$ is the neutral fraction.

Obtaining the ion fraction f_{ion} is quite complicated as ionization depends on the local radiation field and the population of low-energy (primarily 2–10 MeV) cosmic rays that may not follow the same transport as the GeV cosmic rays we model here. Including a more complete treatment of ionization will be the subject of future work, but for now, we remain agnostic about the exact ion fraction for clouds of varying densities and temperatures. We prescribe an ion fraction that, near 10^4 K, matches the tabulated Sutherland & Dopita (1993) ionization fraction fairly well. Note that, while the Sutherland–Dopita table assumes collisional ionization equilibrium and only extends down to 10^4 K, we extend the neutral fraction to lower temperatures following a hyperbolic tangent function. We then impose a floor on the ion fraction, $f_{\text{ion}}^{\text{min}}$, ranging from 10^{-4} to 10^{-1} to parameterize our ignorance and explore the sensitivity of our results to ionization level, although not reaching the f_{ion} characteristic of Galactic molecular clouds, which can be as low as $f_{\text{ion}} \sim 10^{-8}$. For the relatively diffuse clouds we focus on, with column densities $\sim 10^{20} - 10^{21} \text{ cm}^{-2}$, it is likely that there are sufficient UV photons to keep carbon ionized, and ion fractions lower than 10^{-4} are unlikely (Hollenbach & Tielens 1999; Neufeld & Wolfire 2017; Silsbee & Ivlev 2019). The equation we use is

$$f_{\text{ion}} = 1 - \left(\frac{1}{2} (1 - f_{\text{ion}}^{\text{min}}) \left(1 + \tanh \left(\frac{a - T}{c} \right) \right) \right), \quad (8)$$

where T is the temperature in Kelvin, $a = 1.6 \times 10^4$ K, and $c = 2 \times 10^3$ K.

Figure 1 shows the results of Equation (6) for ionization fraction functions with various $f_{\text{ion}}^{\text{min}}$. For a large cosmic-ray energy e_{CR} relative to magnetic energy e_B , the streaming velocity tracks pretty closely to the ion Alfvén velocity as the last term on the right-hand side (RHS) of Equation (6) is small. We also see, however, that the ion Alfvén velocity and streaming velocity can diverge considerably when the minimum ion fraction is higher or when the cosmic-ray energy density is lower. Clearly the streaming velocity is sensitive to ion fraction function, and this is especially true at cloud interface temperatures near 10^4 K. Exploring this sensitivity with more realistic treatments of ionization will be the subject




Figure 1. Top panel: prescribed ion fraction as a function of temperature for three minimum ion values, $f_{\text{ion}}^{\text{min}} = 10^{-4}$, 10^{-2} , and 10^{-1} . Bottom panel: ion Alfvén velocities (v_A^{ion}) and effective streaming velocities ($v_{\text{st}}^{\text{eff}} = v_A^{\text{ion}} + \kappa_{\parallel}/(\gamma_{\text{CR}} I_{\text{CR}})$) from Equation (6) for two different $f_{\text{ion}}^{\text{min}}$ values and varying ratios of magnetic to cosmic-ray energy density. For high cosmic-ray energy density, $v_{\text{st}}^{\text{eff}} \approx v_A^{\text{ion}}$ when the ion fraction is low, but these velocities can diverge significantly when the ion fraction is higher or cosmic-ray energy density is lower. This is especially true near 10^4 K, where ion–neutral damping induces a bump in $v_{\text{st}}^{\text{eff}}$.

of future work, but as we will see in our simulations of cosmic-ray fronts, overall cosmic-ray penetration into clouds is largely insensitive to whether we include or exclude ion–neutral damping (the RHS of Equation (6)). This is because the dynamics are set by the cosmic-ray energy density outside the cloud, where e_{CR} is by design comparable to the thermal and magnetic energy densities in our simulations; this is the case in the average Milky Way ISM and near cosmic-ray sources, and, rather than situations where e_{CR} is small and therefore will not significantly affect the ISM, this is our regime of interest.

Of course, when there is no cosmic-ray pressure gradient, cosmic rays are completely decoupled from waves and free-stream, but when a pressure gradient does exist, cosmic rays generate self-confining waves. We show in our simulations that, in our regime of interest when e_{CR} is not much less than e_B , a steep pressure gradient can develop when the streaming velocity decreases in a small region at the cloud interface (compare $v_{\text{st}}(T = 3 \times 10^4 \text{ K})$ to $v_{\text{st}}(T = 10^5 \text{ K})$, and subsequent transport follows v_A^{ion} fairly closely. The necessary pressure

gradient to overcome ion–neutral damping can even be created from an initially uniform cosmic-ray pressure when collisional losses are strong (see Appendix and Skilling & Strong (1976)).

This is all consistent with the work of Everett & Zweibel (2011), who study the steady-state interaction between a population of GeV-energy cosmic-ray protons, with a nonzero pressure gradient, and cold clouds of varying densities and magnetizations. They find that steeper pressure gradients and higher magnetic field strengths allow cosmic rays to remain coupled to waves deeper into the cloud. For a shallow pressure gradient, though, the cosmic rays eventually decouple from waves and free-stream through the cloud, leading Everett & Zweibel (2011) to hypothesize that the cosmic-ray energy inside and outside clouds will be roughly the same. Whether the diffusive flux can compensate for the drop in advective flux, leading to this scenario, depends on the dynamical and thermodynamical effects of the cosmic-ray pressure gradient, which is accounted for in our simulations but not in Everett & Zweibel (2011), in which the cloud properties are held fixed.

4. Computational Methods

Until recently, a large suite of simulations probing the interaction of cosmic rays with cold, partially neutral ISM clouds would have been infeasible. The heart of the issue is that self-confined cosmic rays can only stream down their pressure gradient. To avoid grid-scale instabilities caused by this abrupt direction change at cosmic-ray extrema, the predominantly employed “regularization” method (Sharma et al. 2009) adds artificial diffusion, which necessitates a quadratic time-step restriction ($\Delta t \propto (\Delta x)^2$). For fast cosmic-ray transport (hence, fast signal speed) in partially neutral gas, this condition makes high-resolution simulations imprudent.

We can overcome this limitation using a new cosmic-ray treatment based on a two-moment method previously used for radiative transfer (Jiang & Oh 2018), which has been implemented in Athena++ (Stone et al. 2020). Unlike the regularization method, this method has no dependence on the smoothing parameter near cosmic-ray maxima and boasts a linear stable time-step scaling ($\Delta t \propto \Delta x$) that makes it computationally tractable to simulate very fast transport speeds and resolve the small-scale structures that induce them.

We refer the reader to Jiang & Oh (2018) for more details (see also a similar method by Thomas & Pfrommer 2019) and use this space just to outline our additions of ionization-dependent transport. Namely, we implement a flag, which, when turned on, calculates the streaming velocity as the ion Alfvén velocity everywhere (including for the $v_A^{\text{ion}} \cdot \nabla P_{\text{CR}}$ heating term), given the temperature-dependent ion fraction function of Equation (8). This flag also includes the additional boost due to ion–neutral friction as a diffusive flux term with diffusion coefficient given by Equation (5).

To properly model streaming with the two-moment method, there is a maximum speed-of-light parameter, V_m , that needs to be much greater than the maximum propagation speed. For simulations with $v_{\text{st}} = v_A$ (assuming a fully ionized gas), we set $V_m = 10^9 \text{ cm s}^{-1}$, which is more than sufficient because the gas Alfvén speed is safely lower than 10^8 cm s^{-1} in all simulations. For simulations including ionization-dependent transport in partially neutral gas, we increase V_m to $10^{10} \text{ cm s}^{-1}$. In this case, v_A^{ion} approaches 10^8 km s^{-1} in our fiducial cloud setup and locally exceeds that in the multicloud simulations of Section 6. To ensure that V_m remains safely higher than the fastest

Table 1
Simulation Parameters for the Single-cloud Simulations of Sections 5.1 and 5.2

r_c, t_c	x_c	$n_{\text{cold}}, T_{\text{cold}}$	$n_{\text{hot}}, T_{\text{hot}}$	$\mathbf{B} = B\hat{x}$	F_{CR}^0
10 pc, 5 pc	1 kpc (1D, Section 5.1), 100 pc (1D+2D, Section 5.2) 100 pc (3D, Section 5.2)	$10 \text{ cm}^{-3}, 10^3 \text{ K}$	$0.1 \text{ cm}^{-3}, 10^5 \text{ K}$	$5 \mu\text{G}$ (1D, Section 5.1), $1\text{--}5, 10 \mu\text{G}$ (1D+2D, Section 5.2) $1.5 \mu\text{G}$ (3D, Section 5.2)	$1.54 \times 10^{-5} \text{ erg cm}^{-2} \text{ s}^{-1}$
Pulse Duration	Grid Size	Resolution	V_m	$\max v_A^{\text{ion}}$	$\max \kappa_{ }$
30 Myr	(2, 0.25, 0.25)kpc	0.5 pc (1D), 1 pc (2D+3D) (see Appendix A)	10^9 cm s^{-1} ($f_{\text{ion}}^{\text{min}} = 1.0$) $10^{10} \text{ cm s}^{-1}$ ($f_{\text{ion}}^{\text{min}} = 10^{-4}$)	10^8 cm s^{-1} ($f_{\text{ion}}^{\text{min}} = 1.0$) 10^9 cm s^{-1} ($f_{\text{ion}}^{\text{min}} = 10^{-4}$)	$3 \times 10^{29} \text{ cm}^2 \text{ s}^{-1}$ ($f_{\text{ion}}^{\text{min}} = 1.0$) $3 \times 10^{30} \text{ cm}^2 \text{ s}^{-1}$ ($f_{\text{ion}}^{\text{min}} = 10^{-4}$)

Note. Given are the cloud radius, r_c , and interface thickness, t_c ; cloud distance, x_c , from the left boundary; cloud number density, n_{cold} , and temperature, T_{cold} ; background number density, n_{hot} , and temperature, T_{hot} ; magnetic field strength, $\mathbf{B}$, which is $5 \mu\text{G}$ in Section 5.1 but varied between 1 and $10 \mu\text{G}$ in Section 5.2; cosmic-ray flux at the left boundary, F_{CR}^0 ; duration of the cosmic-ray pulse; 3D grid dimensions in kiloparsecs (note that only the first one and two dimensions are relevant for 1D and 2D simulations, respectively); maximum resolution (see Appendix A for a convergence study); maximum speed-of-light parameter, V_m , which varies depending on minimum ion fraction, $f_{\text{ion}}^{\text{min}}$, capped streaming speed; and capped diffusivity. Where appropriate, differences between the 1D, 2D, and 3D simulations are listed.

propagation speed, we also cap the streaming speed at 10^9 cm s^{-1} and $\kappa_{||}$ at $3 \times 10^{30} \text{ cm}^2 \text{ s}^{-1}$. We tested the validity of these caps in 1D simulations by increasing them each by a factor of 10 (and therefore increasing V_m by a factor of 10 as well), and we found negligible differences.

A separate issue is whether cosmic rays with high effective streaming speeds can be treated as a fluid. We will revisit this in Appendix B, but for now, we take a detailed look at cosmic-ray–cloud interactions within the fluid assumption.

We also include sink terms $-\Lambda_C$ and $-\Lambda_H$, corresponding to Coulomb and hadronic collisions, respectively, in the cosmic-ray energy equation. Corresponding $+\Lambda_C$ and $+\Lambda_H/6$ terms in the thermal energy equation account for (1) the fact that all cosmic-ray energy lost to Coulomb interactions heats the gas, and (2) on average, one-sixth of the energy lost to hadronic collisions ends up in secondary $e^\pm$ pairs, the majority of which heat the gas through Coulomb collisions (Pfrommer et al. 2017). The rest of the hadronic energy-loss escapes as gamma-rays and neutrinos. We use the equations of Enßlin et al. (2007) and Pfrommer et al. (2017) for the Coulomb and hadronic loss terms:

$$\Lambda_{\text{coul}} = -2.78 \times 10^{-16} \left(\frac{n_e}{\text{cm}^{-3}} \right) \left(\frac{e_c}{\text{erg cm}^{-3}} \right) \text{erg cm}^{-3} \text{s}^{-1}, \quad (9)$$

$$\Lambda_{\text{hadr}} = -7.44 \times 10^{-16} \left(\frac{n}{\text{cm}^{-3}} \right) \left(\frac{e_c}{\text{erg cm}^{-3}} \right) \text{erg cm}^{-3} \text{s}^{-1}, \quad (10)$$

where n is the number density of nucleons, and n_e is the number density of free electrons.⁷

Note that, while the gas evolution equations contain collisional and collisionless heating terms, we do not include radiative cooling. This is a practical choice that allows us to

⁷ During the write-up of this project, we found an error in our code where we calculated the Coulomb loss rate using the total gas density instead of the density modified by ionization fraction. Coulomb losses, which should become negligible in cold, neutral gas, are then overestimated in our simulations. Because hadronic collisions dominate everywhere, this is a small correction to the total collisional loss rate and one that is, in reality, partially offset by our neglect of ionization losses that act primarily in dense gas.

prescribe initially static gas distributions and isolate the already complex interplay with cosmic rays. We discuss the implications of this further in Section 7 and plan to follow up with a future project including cooling.

5. MHD Simulations: Single Cloud

We ran a large number of 1D, 2D, and 3D simulations of different cloud sizes and densities embedded in various environments. For our single-cloud simulations, the parameter choices we present in this paper are given in Table 1, including which subsections they apply to. We will focus here on the interaction between a cosmic-ray front and one representative cloud with maximum total number density $n_{\text{cold}} = 10 \text{ cm}^{-3}$ in a background density of $n_{\text{hot}} = 0.1 \text{ cm}^{-3}$. The initial density profile takes the same form as in Wiener et al. (2017a, 2019):

$$\rho(r) = \rho_{\text{hot}} + \frac{1}{2}(\rho_{\text{cold}} - \rho_{\text{hot}}) \left(1 - \tanh \left(\frac{r - r_c}{t_c} \right) \right). \quad (11)$$

The fiducial cloud radius, r_c , is 10 pc and tapers off with an interface that is fiducially $t_c = 5 \text{ pc}$ thick. As we will see, the interface thickness and how many cells resolve the interface are key parameters. The initial setup is a 2 kpc long box threaded by a constant magnetic field in the x direction (plasma $\beta \approx 1.6$ for a $5 \mu\text{G}$ field) and at constant thermal pressure $3.23 \times 10^{-12} \text{ dyne cm}^{-2}$, giving a cloud temperature of 10^3 K and a background temperature of 10^5 K . We do not include radiative cooling, so such an unstable temperature poses no issue. At a temperature of 10^5 K , the background medium is safely fully ionized, which isolates the onset of super-Alfvénic streaming to the cloud region. The cosmic-ray energy density is set to an initially negligible value.

For the top and bottom grid boundaries ($\hat{y}$ and $\hat{z}$), we use standard outflow boundaries (values are copied over between the ghost cells and edge cells), while we use user-defined boundary conditions along the direction of the magnetic field ($\hat{x}$). Cosmic rays are injected at the left boundary by setting the cosmic-ray energy flux, which follows a time profile that gradually increases and then gradually shuts off after 30 Myr:

$$F_{\text{CR}} = F_{\text{CR}}^0 (1 - e^{-t/3})(1 - e^{-(t-30)/3}), \quad (12)$$

where our fiducial value of F_{CR}^0 is $1.54 \times 10^{-5} \text{ erg cm}^{-2} \text{ s}^{-1}$ and t is in Myr. By considering a finite injection period, we envision a local energy burst from a star cluster, with a characteristic lifetime of order 30 Myr, acting on a cloud in the disk or inner galactic halo. In the outer galactic halo, where the collective effects of supernovae in the disk may be described as an approximately constant, vertical energy flux, a continuous energy injection over a few hundred Myr may also be appropriate. For pure streaming transport, this yields a cosmic-ray front with pressure of order the thermal and magnetic pressures, which in 2D and 3D is significant enough to warp magnetic field lines in the lateral direction when cosmic-ray bottlenecks form at cold cloud interfaces. This is exaggerated further in simulations with lower magnetic field strength or with larger cosmic-ray influxes, as we show in Section 5.2. All other variables abide by outflow boundary conditions at this left boundary. Note that this method of injecting cosmic rays differs from that of Brüggén & Scannapieco (2020), who held the cosmic-ray energy density fixed at the left boundary, but is very similar to Wiener et al. (2019), who implemented a cosmic-ray energy source term at the left boundary.

At the right boundary, all quantities follow outflow boundaries except for the cosmic-ray energy density: we set this value to be slightly smaller in the ghost cells than in the neighboring domain cells to ensure that cosmic rays leave the grid due to the negative pressure gradient. In practice, our right boundary is always far enough away that the Alfvén travel time to the boundary is long compared to the evolution time we consider, and the exact boundary condition does not matter. We confirmed this by running some test simulations with varying fractions of the ghost cell to edge cell cosmic-ray energy densities and some with the boundary placed twice as far away, and we found our results to be insensitive to these choices. As a test, we reran the simulations of Wiener et al. (2019) and exactly matched their results.

As the cosmic-ray front builds up, cosmic rays stream down their pressure gradient toward the cloud at the Alfvén speed ($\approx 29 \text{ km/s}$ for a $5 \mu\text{G}$ field). The sound speed outside the cloud is $\approx 37 \text{ km/s}$, and as found in Wiener et al. (2017a, 2019), the sound wave that outpaces the cosmic-ray front distorts and begins to push the cloud before the cosmic rays arrive. We also ran some tests with a background temperature of $3 \times 10^4 \text{ K}$, thereby decreasing the sound speed below the Alfvén speed, and we found our main conclusions about cosmic-ray–cloud interactions to be robust. Of course, decreasing the temperature while keeping the same cloud density means the cloud is now less pressurized, so there are some differences in the resulting cloud morphology.

5.1. 1D Simulations

We begin with 1D simulations, which therefore force the cosmic rays to eventually enter the cloud. We place the fiducial $n_{\text{cloud}} = 10 \text{ cm}^{-3}$ cloud at the center of the box a distance $x = 1 \text{ kpc}$ from the left boundary. We also ran simulations with the cloud placed closer to the source and find qualitatively similar results. In Figure 2, we focus on a cloud with interface width 5 pc, and in Figure 3, we show the effect of decreasing the interface width to 0.5 pc. In each set of simulations, the resolution is set to 1/10 of the interface width,

i.e., 0.5 pc and 0.05 pc, respectively. All other parameters are kept fixed (see Table 1).

Figure 2 shows a time series of the cosmic-ray front approaching and entering the cloud for a simulation with $f_{\text{ion}}^{\text{min}}$ set to 10^{-4} . The top panel shows the gas density and cosmic-ray energy density, while the bottom panel shows v_A^{ion} and $\kappa_{\parallel}$. Although the cosmic-ray pressure within the cloud is initially constant, spatially dependent collisional losses sustain a small pressure gradient directed toward the cloud center, i.e., $l_{\text{CR}} > 0$. Because the cosmic-ray energy density within the cloud is orders of magnitude less than the thermal and magnetic energy densities, the cosmic-ray pressure gradient is still very small, thereby giving a huge effective diffusivity (see Equation (5)).

Cosmic-ray propagation quickly changes when the cosmic-ray front approaches the cloud. The bottom panel of Figure 2 shows that, in the cloud interface, before the temperature decreases below $\approx 10^4 \text{ K}$, v_A^{ion} decreases considerably. Here, the gas is still mainly ionized, and the increase in density outpaces the decrease in ion fraction. Once the temperature decreases further, the neutral fraction and, hence, v_A^{ion} increase significantly, but not before this decrease in propagation speed has started a cosmic-ray traffic jam. This bottleneck amplifies the preexisting cosmic-ray pressure gradient (leading to small l_{CR}) at the cloud interface, which suppresses the diffusive flux due to ion–neutral damping. The resulting cosmic-ray propagation is then dominated by advection at the ion Alfvén speed, as cosmic rays are locked to the wave frame. This is seen in the bottom panel of Figure 2, where the diffusivity drops from the capped value of $3 \times 10^{30} \text{ cm}^2 \text{ s}^{-1}$ down to $\approx 10^{25} \text{--} 10^{26} \text{ cm}^2 \text{ s}^{-1}$.

We check the effects of ion–neutral damping further by varying $f_{\text{ion}}^{\text{min}}$ between 10^{-4} and 1.0 in Figure 3. Higher $f_{\text{ion}}^{\text{min}}$ values correspond to lower advective fluxes and higher diffusive fluxes, as expected from Section 3 and Figure 1, so we expect the effects of ion–neutral damping to be amplified at higher $f_{\text{ion}}^{\text{min}}$. This appears true in localized regions, but the advective flux still dominantly determines the transport. Compared to $f_{\text{ion}}^{\text{min}} = 10^{-4}$ and 10^{-2} , which have almost identical density and cosmic-ray energy profiles, for $f_{\text{ion}}^{\text{min}} = 10^{-1}$, the diffusivity is boosted, especially at the cloud interface region near $T = 10^4 \text{ K}$. This partially suppresses bottleneck formation at the leading cloud edge (the cloud morphology no longer shows a pronounced bump) and drives a flat cosmic-ray pressure profile in isolated regions near the cloud edges. Cosmic-ray pressure gradients still drive waves throughout the cloud interior, although locking cosmic rays to waves and suppressing the diffusivity. Because v_A^{ion} is now lower, there is an overall slower cosmic-ray passage through the cloud and a cosmic-ray profile intermediate between $f_{\text{ion}}^{\text{min}} = 10^{-2}$ and $f_{\text{ion}}^{\text{min}} = 0$ (fully ionized). Clearly the nonlinear transport is especially complex in this case, but the cosmic-ray energy downstream of the cloud is surprisingly very similar to that with lower $f_{\text{ion}}^{\text{min}}$ values.

We also vary the cloud interface width from 5 pc to 0.5 pc, shown by the dashed lines in Figure 3, keeping the interface resolved by 10 cells (0.05 pc resolution). We find that even this thin interface induces a severe bottleneck, and the amount of cosmic-ray energy leaving the back side of the cloud is comparable to the thicker interface (0.5 pc) case. In light of the limited resolution of most galaxy-scale simulations, it is instructive to point out that degrading the resolution does lead to a larger cosmic-ray flux out the back side of the cloud; in




Figure 2. Time series of a cosmic-ray front impacting an initially stationary, partially neutral cloud of radius 10 pc, interface width of 5 pc, and threaded by a constant $5 \mu\text{G}$ magnetic field. The blue lines show the density, while the black lines show the cosmic-ray energy density. The minimum ion fraction we consider is $f_{\text{ion}}^{\text{min}} = 10^{-4}$, resulting in an increase in the ion Alfvén speed and time-varying diffusivity within the cloud (both shown in the bottom panel). The cloud initially is at 1.0 kpc away from the left boundary, which corresponds to 0 pc in these plot coordinates, but the combination of cosmic-ray pressure and preceding acoustic wave pushes the cloud to the right over time. Note that the rightmost panel is shifted to follow the cloud as it moves ≈ 100 pc over 40 Myr. Even with ionization-dependent transport included, cosmic rays bottleneck at the leading edge, where they encounter a dip in transport speed. The ensuing pressure gradient drives waves and suppresses the diffusivity due to ion-neutral damping. The resulting cosmic-ray pressure gradient is steep at the front and back cloud edges.




Figure 3. Density (blue lines) and cosmic-ray energy density (black lines) for the fiducial cosmic-ray front impinging on a cloud of radius 10 pc with magnetic field strength $5 \mu\text{G}$. We vary the interface width from 5 pc (solid lines) to 0.5 pc (dashed lines), keeping the interface resolved by 10 cells in each case (0.5 pc and 0.05 pc resolution, respectively). The panels correspond to different $f_{\text{ion}}^{\text{min}} = 10^{-4}, 10^{-2}, 10^{-1}$, and 1.0 ($v_{\text{st}} = v_A$) going from left to right.

these cases, what should be steep cosmic-ray pressure gradients at both the front and back interfaces are smoothed out, leading to an enhancement of the diffusive flux. We expand on this further in the Appendix A.

We ran the same simulation assuming the medium is fully ionized and compare results in Figure 4. The top row shows both collisional and collisionless energy loss. As expected, and as seen in Figure 2, the steep pressure gradient that forms at the front of the cloud collisionlessly transfers cosmic-ray energy to waves at a rate $\propto v_A^{\text{ion}} \cdot \nabla P_{\text{CR}}$. Wave damping then heats the

cloud edge, as seen in the bottom panel of Figure 4, and broadens the interface, resulting in a thumb-like density compression that moves upstream.

In the fully ionized case, this heating tracks the cosmic-ray front, which, at this time stamp, has slowly moved to the middle of the cloud. In the partially neutral case, fast transport quickly flattens the cosmic-ray pressure in the cloud, so heating is negligible in the cloud center. When cosmic rays hit the back of the cloud, though, they encounter a drop in streaming speed as the medium transitions from neutral to ionized (high v_A^{ion} to




Figure 4. Fiducial cosmic-ray front impinging on a cloud of radius 10 pc and interface width 5 pc with $f_{\text{ion}}^{\text{min}} = 10^{-4}$ (left) and $v_{\text{st}} = v_A$ (right). The top row shows energy losses (both collisional and collisionless). The bottom row shows density and temperature.

low v_A^{ion}). This second bottleneck persists for a long time while cosmic-ray energy slowly increases downstream of the cloud, and the steep pressure gradient there induces collisionless energy loss, just as it does on the front cloud edge. This heating may have interesting implications for ion abundances and kinematics (Wiener et al. 2017a); however, we note that the energy-loss rate is only of order $10^{-26} \text{ erg cm}^{-3} \text{ s}^{-1}$, which, for a density of $1\text{--}10 \text{ cm}^{-3}$, is far less than the expected cooling rate of order $10^{-23} \text{--} 10^{-26} \text{ erg cm}^{-3} \text{ s}^{-1}$ or higher at temperatures of a few $\times 10^4 \text{ K}$. In future work, we will include radiative cooling to address this directly and determine realistic interface widths, which are set by a balance between cooling, conduction, and cosmic-ray heating. Whether low-energy cosmic rays also modify this environment by increasing the ionization fraction is an open question similarly beyond the scope of this paper.

5.2. 2D and 3D Simulations

To explore the effects of higher dimensions and build intuition for our main suite of simulations in Section 6, we study the same cosmic-ray front interacting with a 2D cloud and, in two of our simulations with $B = 1 \mu\text{G}$ and $5 \mu\text{G}$, with a 3D cloud. In the cases we will show, the cloud density (10 cm^{-3}), radius (10 pc), and interface width (5 pc) are the same as in the 1D simulations (Section 5.1), with an initially constant magnetic field in the x direction, but we move the cloud to a closer distance of 100 pc from the boundary. This is more akin to a cosmic-ray front impacting dense clouds within an ISM scale height of a few hundred parsecs. To make a fair comparison between 1D, 2D, and 3D, we reran our 1D simulations with the cloud at a distance of 100 pc instead of 1 kpc. The 2D and 3D simulation resolution is $\approx 1 \text{ pc}$ out to a distance of 1 kpc, where the grid coarsens to a resolution of 16 pc. As in the 1D simulations, the left boundary is placed 2 kpc away, and here we focus our analysis on the closest 1 kpc, where the cloud is most highly resolved.

Figure 5 shows snapshots from a subset of our 2D simulations with magnetic field strengths of $1 \mu\text{G}$ (bottom row) and $5 \mu\text{G}$ (top row). The cloud properties are the same as in Section 5.1, but the extra dimensionality means that not all field lines (and hence, cosmic rays) enter the cloud, and field




Figure 5. Slices of cosmic-ray energy density for various 2D, single-cloud simulations. The cloud density initially peaks at 10 cm^{-3} , and the white and yellow contours represent densities of 1 and 10 cm^{-3} , respectively. Varying the magnetic field strength from $5 \mu\text{G}$ (top row) to $1 \mu\text{G}$ (bottom row), we see distinct differences in magnetic field topology created by the pressure wave that precedes the cosmic-ray front. In the fully ionized case (left), cosmic-ray propagation traffic-jams at the front cloud edge, casting a shadow in E_{CR} downstream, while allowing for ionization-dependent transport through the cloud (right) leads to a slightly less disrupted cloud, greater E_{CR} downstream, and a second bottleneck at the back edge seen most clearly in the lower-right panel. In the upper-right panel, there are still bottlenecks at the front and back cloud edges, but the cosmic-ray pressure drops are less noticeable, reflecting the faster transport that allows cosmic rays to more easily fill into the cloud along straight magnetic field lines.

lines are allowed to evolve. For varying magnetic field strengths (top versus bottom rows) and varying ion fraction (left versus right columns), we see differences in how cosmic rays penetrate and pressurize the cloud and in the cloud morphology (shown by white and yellow contours of density 1 and 10 cm^{-3} , respectively).

In the $B = 5 \mu\text{G}$ cases, magnetic field line warping is negligible due to strong magnetic tension, and cosmic rays propagate straight across the cloud. Assuming full ionization, v_{st} decreases sharply in the cloud, inducing a bottleneck on the front edge and a cosmic-ray shadow behind the cloud. This shadow fills in very slowly because the collisional loss time in the cloud is significant compared to the transport time. The cloud morphology is very similar to what is seen in Wiener et al. (2019) and Brügger & Scannapieco (2020), where the top and bottom cloud edges with the lowest gas column get pushed and form tails behind the cloud, followed eventually by full-cloud acceleration to the right. When ionization-dependent transport is included, the formation of these tails is not as dramatic because the cosmic-ray pressure profile becomes flatter in the cloud and therefore does not push on the gas as much. The cosmic-ray pressure behind the cloud is clearly larger, as well, because the cosmic-ray transport time is shorter. A second bottleneck, as seen in the 1D simulations, is faintly visible on the back side of the cloud where the transport speeds drop precipitously to the gas Alfvén speed, v_A instead of v_A^{ion} .

With a magnetic field strength of $B = 1 \mu\text{G}$, corresponding to a plasma beta of 75 instead of 3, the magnetic field topology plays a more significant role. Before the cosmic-ray front even reaches the cloud, the preceding acoustic wave disturbs the magnetic field and warps it around the cloud. When cosmic rays approach the cloud, they must follow the stretched field




Figure 6. Profiles at $t = 60$ Myr down the cloud axis showing the evolution of density and cosmic-ray energy density for the fiducial cloud in 1D (solid lines), 2D (dashed lines), and for two simulations with $B = 1 \mu\text{G}$ and $5 \mu\text{G}$, 3D (dotted-dashed lines). Each panel is for a different magnetic field strength, and each simulation is without ionization-dependent transport ($f_{\text{ion}}^{\text{min}} = 1.0$). The cloud is initially placed 100 pc from the cosmic-ray source but accelerates to varying degrees depending on magnetic field strength and dimensionality. In 1D, the cloud acceleration monotonically decreases with increasing magnetic field strength, but in 2D, there is a sweet spot around $B = 3\text{--}5 \mu\text{G}$ where the cosmic-ray front is most effective at pushing the cloud.




Figure 7. Same as Figure 6 but with $f_{\text{ion}}^{\text{min}} = 10^{-4}$.

lines that are draping around the cloud and therefore take longer to enter. For field lines that do penetrate the cloud, regardless of transport, the slower Alfvén speed outside the cloud leads to a buildup of cosmic-ray pressure and a severe bottleneck on the front edge.

5.2.1. Cloud Acceleration: 1D versus 2D versus 3D

In each simulation, the cloud is eventually accelerated but to different extents depending on transport treatment, magnetic field strength, and dimensionality. Figures 6 and 7 compare

profiles down the cloud midplane of cosmic-ray energy density and gas density after $t=60$ Myr. Solid lines show the 1D simulations, dashed lines show the 2D simulations, and in two cases, we ran parsec-resolution 3D simulations without ionization-dependent transport, shown by the dotted-dashed lines.

Let us first analyze the 1D results without ionization-dependent transport (the solid lines in Figure 6), which are the most straightforward to interpret. Cloud acceleration depends on the magnitude and duration of the cosmic-ray pressure gradient. Keeping in mind that the cosmic-ray pulse only lasts for a finite 30 Myr, at which point the upstream pressure will fade at a rate proportional to the Alfvén speed, the cloud gets a push only until the cosmic-ray pressures upstream and downstream equilibrate. At late times, when the pressure has equilibrated, the cosmic-ray energy density slowly drops due to the collisional energy sink inside the cloud and also cosmic rays escape out the right box edge, which occurs because we enforce a cosmic-ray gradient at that boundary. For a higher Alfvén speed (higher magnetic field strength), cosmic rays build up a short-lasting and relatively low pressure upstream, and because they travel more quickly through the cloud, they lose less energy to collisions and emerge with a no-negligible pressure downstream. So, the cosmic-ray pressure gradient is small and also short lived, leading to the least cloud acceleration for the $B=10\ \mu\text{G}$ simulation. As seen in Figure 6, the distance the cloud travels monotonically increases for decreasing field strength.

In 2D, the magnetic field topology is a complicating factor, as we see the low and high magnetic field cases both give very little cloud acceleration. For high magnetic field strength, field line warping is negligible, and the setup is effectively 1D. Again, the pressure gradient is relatively small, and the pressures equilibrate quickly to give very little cloud acceleration. On the contrary, for low magnetic field strength, the field lines warp so much that it allows cosmic rays to move around the cloud, taking pressure from the upstream, while cosmic rays instead converge and build up behind the cloud. The cosmic-ray pressure drop across the cloud is small, which again leads to relatively little cloud acceleration.

Intermediate magnetic field strengths $B \approx 3\text{--}5\ \mu\text{G}$ ($\beta \approx 4.5\text{--}1.6$) give a sweet spot for cloud acceleration. In these cases, field lines do not warp as much, so the upstream pressure is higher, while the cosmic-ray transit time through the cloud is intermediate (of order 5–10 Myr and comparable to the collisional loss time), so the downstream pressure is low. This creates a sizable and fairly long-lasting pressure gradient; however, note that the 3D simulations show flat cosmic-ray pressure profiles after 60 Myr and clouds that have better maintained their shape and moved less distance. While the differences between 2D and 3D are small compared to the differences between 1D and 2D, clearly the extra dimensionality plays a role, motivating a larger sample of 3D simulations that we will run in the future.

5.2.2. Effects of Ionization-dependent Transport

Figure 7 shows the same 1D versus 2D comparison as before, but we now turn on ionization-dependent transport with $f_{\text{ion}}^{\text{min}}=10^{-4}$. Again, for the 1D simulations, the cloud distance increases monotonically for decreasing magnetic field strength, as the same physical picture of pressure equilibration applies. Note, however, that for a given magnetic field strength, the

fully ionized case always gives a slightly greater cloud acceleration than when partially neutral gas is accounted for. This difference is most dramatic when the magnetic field strength is high and ionization-dependent transport through the cloud leads to an especially quick pressure equilibration; for low magnetic field strengths, the bottleneck at the cloud interface is severe and results in a pressure gradient that pushes the cloud to nearly the same extent regardless of $f_{\text{ion}}^{\text{min}}$. While the cosmic-ray profiles are markedly different inside the clouds (steep for $f_{\text{ion}}^{\text{min}}=1.0$ and flat for $f_{\text{ion}}^{\text{min}}=10^{-4}$), the total pressure drop from upstream to downstream is comparably large in both cases, and it is the total cosmic-ray pressure difference across the cloud, regardless of the smaller-scale pressure profile differences, that determines the force.

In 2D, efficient cloud acceleration again disfavors high magnetic field strengths, but clearly, ionization-dependent transport has suppressed cloud acceleration in the intermediate field regime as well, as the density peaks shift to at most 200 pc after 60 Myr. In the fully ionized case, intermediate field strengths led to a sweet spot with only small field line warping and moderate transport speed across the cloud; combined, this created a large pressure difference. Ionization-dependent transport, however, equilibrates the pressure on a shorter timescale, breaking down the sweet spot. Low magnetic field strengths allow for a greater buildup of cosmic-ray pressure upstream, but cosmic-ray flow around the cloud once again makes the pressure drop across the cloud almost negligible.

Consistent with the morphological differences in Figure 5, ionization-dependent transport keeps the cloud fairly compact, because there is no pressure gradient within the cloud, whereas the $f_{\text{ion}}^{\text{min}}=1.0$ density profiles in Figure 6 show the clouds being stretched significantly, which decreases the average cloud density. While none of our ionization-dependent transport simulations resulted in high cloud velocities, it is intriguing to wonder whether ionization-dependent transport and a larger cosmic-ray flux could lead to a significant acceleration of molecular gas while keeping it intact instead of stretched apart. This will be addressed with future simulations.

Our main conclusion is that cosmic-ray acceleration of partially neutral clouds is less efficient than acceleration of fully ionized clouds, but the differences in imparted momentum are small (less than a factor of 2). This is quantified in Figures 8 and 9, which show the cloud momentum, velocity, and change in mass, defining the cloud to be all gas with temperature $T < 2 \times 10^4$ K. As expected, the momentum and velocities with ionization-dependent transport included are lower by a factor of ≈ 2 compared to the fully ionized case. The cloud masses decrease over time in all cases, and they decrease most quickly for higher magnetic field strengths (higher Alfvén speeds) that most quickly transform the cloud edges into long filaments that are sheared apart by the Kelvin–Helmholtz instability, consistent with findings from Brüggén & Scannapieco (2020). For similar reasons, the mass decreases more for clouds that reach greater velocities, as the cloud-crushing time $t_{\text{cc}} \sim \xi^{1/2} r_{\text{cl}} / v$ is then shorter, where ξ is the cloud overdensity relative to the background, r_{cl} is the cloud radius, and v is the velocity of the ambient medium relative to the cloud (Klein et al. 1994). It is worth noting again that there is no radiative cooling in our simulations, which may suppress the ablation of gas at the cloud boundary and even facilitate cloud growth as the surrounding gas mixes and condenses onto the cloud




Figure 8. For the fiducial 2D cloud, placed 100 pc from the left boundary, these panels show the total 2D momentum, volume-averaged velocity, and change in cloud mass M_{cl}/M_{cl}^0 over time, defining the cloud as all gas with a temperature $T < 2 \times 10^4$ K. For each simulation, we vary the magnetic field strength between 1 and 10 μ G, and $f_{ion}^{min} = 1.0$ in all cases. The greatest cloud acceleration occurs when the magnetic field is intermediate strength $\approx 5 \mu$ G, in which case the field line warping is small, but the Alfvén crossing time between the source and the cloud, which partially determines the duration of the cosmic-ray pressure gradient, is longer than the 10 μ G simulation. Lower magnetic field strengths yield more field line warping around the cloud, which inhibits acceleration. In all cases, the cloud mass decreases over time, which is to be expected because we do not include radiative cooling.




Figure 9. Same as Figure 8 but with $f_{ion}^{min} = 10^{-4}$. Cloud acceleration is decreased by a factor of ≈ 2 when ionization-dependent transport is included, and the greatest acceleration now favors somewhat lower magnetic field strengths.

(Gronke & Oh 2018, 2020; Banda-Barragán et al. 2020; Kanjilal et al. 2021).

6. A Multiphase Cosmic-Ray Obstacle Course

We now show a suite of simulations of cosmic-ray propagation through multiclump “obstacle courses.” Most simulations were run in 2D, as that was the most computationally feasible way to scan our parameter space. Our analysis in Sections 6.1 and 6.2 focuses entirely on the 2D simulations; however, we do run a subset of 3D simulations and quantify their cosmic-ray calorimetry in Section 6.3. We intend to run more 3D simulations in the future to probe differences between 2D and 3D further.

Compared to the single-cloud simulations of Section 5.2, where the cosmic-ray front moved horizontally from left to right, we now rotate the box to help the reader envision the scenario we probe: a cosmic-ray front vertically escaping from a clumpy disk. Table 2 contains the simulation parameters, grid sizes, and resolutions we explore. Table 3 lists explicitly which simulations we ran in 2D and 3D for different clump distributions, magnetic field strengths, and transport models.

The mean density at the bottom boundary in each simulation is $\bar{n} \approx 1 \text{ cm}^{-3}$ and initially falls off exponentially with a scale height $H = 0.25 \text{ kpc}$. The magnetic field is constant in the y direction (vertical direction) with a value of either 1 or 5 μ G, and the pressure is constant throughout the simulation box at $3.23 \times 10^{-12} \text{ dyne cm}^{-2}$. In each simulation, the domain extends to the $y = 4 \text{ kpc}$ top boundary, which is sufficiently

far away from the clumpy ISM to alleviate worries about boundary effects. As in Section 5.2, the boundary conditions perpendicular to the magnetic field direction are set to outflow for all quantities, the bottom boundary injects cosmic rays by setting the cosmic-ray flux (again with $F_{CR}^0 = 1.54 \times 10^{-5} \text{ erg cm}^{-2} \text{ s}^{-1}$ and following the time profile of Equation (12)), and the top boundary is an outflow boundary except for the cosmic-ray energy density, where we again impose the ghost value to be slightly smaller than the domain value to ensure that cosmic rays leave the box.

On top of this tapered density profile, we impose a lognormal distribution of density perturbations (Brüggen 2013) with 20 Fourier components in each of the x , y , and (if in 3D) z directions. As with the mean density, the perturbation amplitude also decreases exponentially with scale height 0.25 kpc in order for the density variance to attenuate above the disk into a smooth halo. The density profile is then

$$n = \bar{n} e^{-y/H} e^{\alpha f(x,y,z)}, \quad (13)$$

$$f(x, y, z) = e^{-y/H} (1 - e^{-y}) \sum_{k_x, k_y, k_z} A_i \sin(2\pi k_x x/L + \phi_i) \sin(2\pi k_y y/L + \theta_i) \sin(2\pi k_z z/L + \eta_i), \quad (14)$$

where x, y, z are in units of kiloparsecs, $H = 0.25 \text{ kpc}$, and $k_x, k_y, k_z \in [10, 30]$. A_i and ϕ_i, θ_i , and η_i are randomly chosen from a uniform distribution between $[0, 1]$ and $[0, 2\pi]$, respectively. The parameter L controls the characteristic size of

Table 2
Simulation Parameters for the Multiple Clump Simulations of Section 6

$\bar{n}, \bar{T}$	Scale Height (H)	$B = B\hat{y}$	F_{CR}^0	Pulse Duration
$1 \text{ cm}^{-3}, 10^4 \text{ K}$	250 pc	$1.5 \mu\text{G}$	$1.54 \times 10^{-5} \text{ erg cm}^{-2} \text{ s}^{-1}$	30 Myr
Grid Size		Resolution		
$2 \times 4 \text{ kpc}$ (2D; $L = 5$)	$1 \times 4 \text{ kpc}$ (2D; $L = 2$)	4 pc (2D; $L = 5$)		
$2 \times 4 \times 1 \text{ kpc}$ (3D; $L = 5$)	$1 \times 4 \times 0.5 \text{ kpc}$ (3D; $L = 2$)	4 pc (3D; $L = 5$)		
		2 pc (2D; $L = 2$)		
		2 pc (3D; $L = 2$)		

Note. Given are the mean number density, $\bar{n}$, and mean temperature, $\bar{T}$, at the lower ($y = 0$) boundary; the scale height H over which the density drops but the gas pressure is initially constant at $P_g = 3.23 \times 10^{-12} \text{ dyne cm}^{-2}$; magnetic field strength, B , which is varied between 1 and $5 \mu\text{G}$; cosmic-ray flux at the bottom boundary, F_{CR}^0 ; duration of the cosmic-ray pulse, 2D and 3D grid dimensions in kiloparsecs, which depend on the clump sizes controlled by the parameter L ; and the resolution, which depends on the parameter L . The maximum speed of light parameter, V_m , capped streaming speed, and capped diffusivity are the same as in Section 5.

Table 3
Multiple Clump Simulations Run (Section 6)

2D	L	α	B (μG)	Transport
	5	1.5	5	$f_{\text{ion}}^{\text{min}} = 1.0, f_{\text{ion}}^{\text{min}} = 10^{-4}$, $f_{\text{ion}}^{\text{min}} = 1.0 + \kappa_{ } = 3 \times 10^{27}$, $f_{\text{ion}}^{\text{min}} = 1.0 + \kappa_{ } = 3 \times 10^{28}$, $f_{\text{ion}}^{\text{min}} = 10^{-4} + \kappa_{ } = 3 \times 10^{27}$
	5	1.5	1	$f_{\text{ion}}^{\text{min}} = 1.0, f_{\text{ion}}^{\text{min}} = 10^{-4}$
	5	1.0	5	$f_{\text{ion}}^{\text{min}} = 1.0, f_{\text{ion}}^{\text{min}} = 10^{-4}$
	5	0.5	5	$f_{\text{ion}}^{\text{min}} = 1.0, f_{\text{ion}}^{\text{min}} = 10^{-4}$
	2	1.5	5	$f_{\text{ion}}^{\text{min}} = 1.0, f_{\text{ion}}^{\text{min}} = 10^{-4}$
	2	1.5	1	$f_{\text{ion}}^{\text{min}} = 1.0, f_{\text{ion}}^{\text{min}} = 10^{-4}$
3D	L	α	B (μG)	Transport
	5	1.5	5	$f_{\text{ion}}^{\text{min}} = 1.0, f_{\text{ion}}^{\text{min}} = 10^{-4}$
	2	1.5	5	$f_{\text{ion}}^{\text{min}} = 1.0$

Note. Most simulations for a given (L, α, B) are run with $f_{\text{ion}}^{\text{min}} = 1.0$ and again with $f_{\text{ion}}^{\text{min}} = 10^{-4}$. For a subset of the 2D simulations, additional parallel diffusion is also included.

perturbations, while α sets the density contrast. The extra factor of $(1 - e^{-y})$ in the perturbation equation excludes the clumps away from the bottom boundary. We found in test simulations that if a dense clump is on the boundary, the outflow boundary we impose will copy the gas density into the ghost cells and therefore draw in a large amount of artificial mass. Because, as we will see, cosmic rays push on clumps differently depending on the transport model, this results in different total masses in the simulation box, which muddies our analysis of gamma-ray emission, gas momentum, etc.

Our clump setup varies, with (L, α) between $(5, 1.5)$ and $(2, 0.5)$, giving different density probability distribution functions (PDFs) shown in Figure 10. For the gas pressure we initialize, at the mean density $\sim 1 \text{ cm}^{-3}$, the temperature is $\sim 10^4 \text{ K}$. For $L = 5$, this gives a porous ISM dominated by large (≈ 50 – 100 pc) over/underdensities. The simulation domain extends to $\pm 1 \text{ kpc}$ in the x direction (horizontal direction) in this case. In simulations with $L = 2$, the box only extends to $\pm 500 \text{ pc}$ in the x direction. With small density structures (≈ 10 – 50 pc) more typical of real ISM clouds, these domain widths are sufficient to present cosmic-ray interactions with clouds of varying densities; however, they necessitate higher resolution to more properly resolve the cloud interfaces, which we have found to be important (Section 5.1). For $L = (5, 2)$ the




Figure 10. Density probability distribution function for each combination of (L, α) we present.

resolution is uniformly $(4, 2) \text{ pc}$ until $y = 2 \text{ kpc}$, at which point we make use of static mesh refinement to decrease the resolution to $(16, 8) \text{ pc}$ and save computation time.

It is important to note that these simulations do not include cooling or gravity and are not meant to substitute for more realistic, stratified ISM simulations, though they do replicate the propagation of cosmic rays through layers of a fixed thickness and mean column density. Having a constant magnetic field strength out to many scale heights is also not realistic, so we focus our analysis on the first kiloparsec above the midplane.

6.1. Qualitative Differences

Figure 11 shows snapshots at $t = 40 \text{ Myr}$ of our 2D, $B = 5 \mu\text{G}$ simulations with mean density $= 1.0 \text{ cm}^{-3}$, $(L, \alpha) = (5, 1.5)$, and $f_{\text{ion}}^{\text{min}} = 1.0$ (left) and $f_{\text{ion}}^{\text{min}} = 10^{-4}$ (right). Figure 12 shows the same setups but with $B = 1 \mu\text{G}$. The plots are oriented such that cosmic-ray flux is input at the lower boundary, and magnetic field lines are vertical. The contours show densities of $n = 1 \text{ cm}^{-3}$ and $n = 10 \text{ cm}^{-3}$. The underdense, intercloud region extends down to densities of $n \approx 10^{-3} \text{ cm}^{-3}$ (see Figure 10). Such a configuration, if entirely ionized, gives a range of Alfvén velocities predominantly between 10^5 and 10^7 cm s^{-1} , but when neutral particles are accounted for,




Figure 11. Snapshots at $t = 40$ Myr for the 2D ISM setup with $B = 5 \mu\text{G}$, mean density $= 1.0 \text{ cm}^{-3}$, and $(L, \alpha) = (5, 1.5)$. Simulation parameters are listed in Table 2. Left panels: fully ionized assumption ($f_{\text{ion}}^{\text{min}} = 1.0$). Right panels: ionization-dependent transport included ($f_{\text{ion}}^{\text{min}} = 10^{-4}$). The top panel shows cosmic-ray energy density and density contours of 1 and 10 cm^{-3} , which shows clear variations in how cosmic rays preferentially penetrate or flow around cold clouds. The second row shows v_A^{ion} , followed by gas temperature, collisionless energy-loss rate ($(v_A^{\text{ion}} \cdot \nabla P_{\text{CR}})$), and collisional loss rate ($\propto e_{\text{CR}} n_{\text{gas}}$), again with dashed contours showing densities of 1 and 10 cm^{-3} . Collisionless heating at the cloud interfaces is apparent in both transport cases, while collisional energy loss is preferentially higher in dense gas when ionization-dependent transport is included. An animated version of this figure, showing the evolution of these panels from $t = 1$ –100 Myr, is available in the HTML version of this article.

(An animation of this figure is available.)

the lowest Alfvén velocity is $\approx 10^6 \text{ cm s}^{-1}$, with ion Alfvén velocities in cloud interiors reaching beyond 10^8 cm s^{-1} .

This difference in propagation speed is clearly reflected in the top row of Figure 11, which shows that cosmic rays with ionization-dependent transport have propagated farther into the

domain and have reached the inner “halo,” where the average density has dropped. Cosmic rays from here will free-stream out of the simulation domain at high Alfvén velocities because $v_A \propto \rho^{-1/2}$ for constant B . In the fully ionized case, the cosmic rays that have managed to propagate farthest are those that have




Figure 12. Same as Figure 11 but with $B = 1 \mu\text{G}$ instead of $5 \mu\text{G}$. Magnetic field lines now warp around the cold clumps, squeezing cosmic rays through the gaps in the ISM instead of funneling them through clouds. This effect dominates over differences in transport, as the left and right columns (varying $f_{\text{ion}}^{\text{min}}$) now show smaller differences. Also note that, compared to the $B = 5 \mu\text{G}$ simulations shown in Figure 11, there is now a larger buildup of cosmic-ray pressure that more effectively accelerates the gas column. This is especially true for the hot gas, which gets pushed out by the cosmic rays sweeping through the underdense channels. An animated version of this figure, showing the evolution of these panels from $t = 1\text{--}100 \text{ Myr}$, is available in the HTML version of this article.

(An animation of this figure is available.)

avoided cold clouds and squeaked through the cracks. As in the 2D, single-cloud simulations, for an initial plasma $\beta = 1.6$, magnetic field line warping plays only a limited role, meaning most cosmic rays are eventually forced into the clouds (and additionally, accelerate the clouds). When accounting for

ionization-dependent transport, cosmic rays leak into the dense cores and quickly reappear on the other side. As in Figure 5, there are sharp steps in cosmic-ray energy at cloud edges, both at the front and the back edges in the $f_{\text{ion}}^{\text{min}} = 10^{-4}$ case, as each interface induces a bottleneck.

For a magnetic field strength of $1 \mu\text{G}$ (Figure 12), magnetic field line warping dominantly determines how cosmic rays navigate the medium. Now, cosmic rays with either transport model reside primarily in the intercloud regions. The lower magnetic field strength also leads to a greater buildup of cosmic-ray pressure that exerts a force on the medium. As in the single-cloud simulations of Section 5.2, such a lower field strength leads to only a small acceleration of the cold gas, but we see now that the hot gas is very efficiently pushed out by the cosmic rays that are forced to squeeze through the underdense channels of the ISM.

6.2. Influence on Clouds

Farber et al. (2018) probed the implications of cosmic-ray decoupling for galactic wind driving and star formation feedback. They compared three transport models: cosmic-ray advection, field-aligned diffusion with a single diffusion coefficient (one- κ model), and field-aligned diffusion with a higher diffusion coefficient in cold gas (two- κ model) to mimic the effects of ion-neutral damping. Comparing the one- κ and two- κ models, they found that boosting the diffusivity in the cold phase led to a faster, hotter outflow because the cold phase received less impulse from cosmic-ray fronts. Changing $f_{\text{ion}}^{\text{min}}$, we find that this trend partially occurs. Figure 13 shows the momentum in the full simulation box as a function of time for our $(L, \alpha) = (5, 1.5)$ ISM setup. The top panel shows the momentum regardless of gas temperature, while the bottom two panels split the momentum into “cold” gas ($T < 2 \times 10^4 \text{ K}$) and “hot” gas ($T > 2 \times 10^4 \text{ K}$). When ionization-dependent transport is included, slightly less momentum is imparted to both the cold and hot gas, but only on the level of tens of percent. As argued in Section 5.2, we attribute this to the bottleneck and subsequent suppression of the diffusive flux at the cloud boundary, where large cosmic-ray pressure gradients still drive waves and enact a force on the cloud.

Regardless of whether ionization-dependent transport is included, the momentum in the cold phase exceeds the momentum in the hot phase, which is consistent with more complete simulations of supernova-driven outflows where mass-loading in the cold phase exceeds that in the hot phase (e.g., Kim et al. 2020). The time profiles of cold and hot momentum are noticeably different, though. The hot gas is quickly swept out by the preceding acoustic pulse and cosmic-ray front that moves through. A short time after the cosmic-ray source turns off at $t = 30 \text{ Myr}$, the cosmic-ray pressure gradient diminishes, and the momentum slightly drops. The cold gas momentum rises a bit more slowly; cosmic-ray bottlenecks take time to build up (on the order of the Alfvén crossing time between the cloud and the source), and the cosmic-ray pressure releases only after a long time because the cold clouds move relatively slowly. Pressure gradients on the cold gas then build to greater amplitudes and longer durations than those on the hot gas, leading to a larger but more gradual acceleration of the clouds.

Rather than ionization-dependent transport leading to changes in cloud acceleration, the largest changes in momentum correlate with changing magnetic field strength. For the same injected cosmic-ray flux, a lower magnetic field strength (slower transport) translates to a higher cosmic-ray pressure gradient to drive the gas column outwards. Both the $f_{\text{ion}}^{\text{min}} = 1.0$ and $f_{\text{ion}}^{\text{min}} = 10^{-4}$ models show an increase in impulse. The cold gas momentum increases because cosmic-ray bottlenecks are




Figure 13. Momentum measured in 2D clump simulations with varying L and α and varying transport, with $f_{\text{ion}}^{\text{min}} = 1.0$ (solid lines) and $f_{\text{ion}}^{\text{min}} = 10^{-4}$ (dashed lines). The top panel shows the momentum just in the cold gas ($T < 8 \times 10^3 \text{ K}$), and the bottom panel shows the momentum in the hot gas ($T > 2 \times 10^4 \text{ K}$). Note the change in y-axis reflecting that most of the momentum, in either transport model, is in the cold phase, but lower magnetic field strengths do increase the hot gas momentum quite substantially as cosmic rays squeeze out the gas between dense clouds. With a lower field strength, the cold momentum increases somewhat, as well, due to an increased cosmic-ray pressure gradient as we saw in Section 5.2. When neutral particles are accounted for, the momentum is lower, but only slightly, compared to the fully ionized case.

severe in either case at the cloud edges. The hot gas momentum increases because cosmic rays squeeze between the gaps in clouds, sweeping out the hot gas in front of them (see Figure 12).

Interestingly, changes in momentum also correlate with the characteristic size of clumps controlled by our L parameter. More momentum is imparted especially to the cold gas when L is smaller; this makes sense as the limit $L \rightarrow 0$ is effectively a 1D simulation, which we saw in Section 5.2, was the most effective at accelerating cold clouds.

6.3. Quantifying Cosmic-Ray Propagation and Energy Loss

While shown qualitatively in Figures 11 and 12, how cosmic rays sample the ISM is shown more quantitatively in Figure 14. The left column shows PDFs of cosmic-ray energy density (top) and gamma-ray emission (bottom), and the right column shows the concentration of cosmic-ray energy and gamma-ray




Figure 14. All panels are for 2D simulations. The left column shows the probability distribution function (PDF) of the cosmic-ray energy density E_{CR} (top) and gamma-ray emission $\sim E_{\text{CR}} n_{\text{gas}}$ (bottom) in 32 temperature bins ranging from $T = 10^3 - 5 \times 10^5$ K. The right column shows the concentrations, obtained by dividing the PDFs in the left column by the temperature PDF, which give a comparison of how concentrated the quantities are compared to the fraction of the ISM at each temperature bin. If the fraction of energy or emission equals the volume fraction at that temperature bin, the concentration will be 1.0. Clearly there are differences depending on the transport model, with $f_{\text{ion}}^{\text{min}} = 1.0$ shown by solid lines and $f_{\text{ion}}^{\text{min}} = 10^{-4}$ shown by dashed lines, as well as varying ISM parameters and magnetic field strengths (different colors). Cosmic rays in all cases preferentially sample the mean density and temperature $\approx 2 \times 10^4$, where the gas is fully ionized, Alfvén speeds are low, and cosmic-ray residence times are long. Ionization-dependent transport leads to preferentially more occupancy in dense, cold clouds but less occupancy at transition temperatures $T \sim 8 \times 10^3$ K. This trade-off leads to roughly the same total gamma-ray luminosity, regardless of transport (Section 6.4). With a lower magnetic field strength, field line warping allows cosmic rays to squeeze through the underdense channels of the ISM, pushing out high-temperature gas in the process. In those cases, cosmic rays are the most concentrated at the mean temperature.

emission relative to the fraction of the volume occupied by each of the 32 temperature bins. We show this concentration because the temperature PDF changes depending on transport mode, magnetic field strength, and ISM parameters. So the clearest comparison between simulations is to divide the actual PDF by the temperature PDF; in this case, a concentration of 1.0 means that the fraction of cosmic-ray energy or gamma-ray emission in a given temperature bin equals the volume fraction within that temperature bin. We restrict our analysis to the clumpiest part of the simulation, which is within 500 pc of the lower boundary.

In all simulations, most cosmic-ray energy resides near the mean density and temperature, partially because that represents a maximum in the density PDF (Figure 10) and partially because that is close to the minimum in v_A^{ion} . This is consistent with previous work that found, using a Monte Carlo method, that cosmic rays in M82-like starburst galaxies would generally sample the mean density (Boettcher et al. 2013). At temperatures below $T \approx 10^4$ K, as cosmic rays encounter clumps, the PDFs and concentrations start to diverge. With ionization-dependent transport (dashed lines), there is a steep

drop in cosmic-ray occupancy near temperatures $T \approx 8 \times 10^3$ K, where v_A^{ion} rapidly increases. At similar temperatures, cosmic rays with $f_{\text{ion}}^{\text{min}} = 1.0$ are actually overabundant. These differences make sense if we imagine a single-cloud-cosmic-ray front interaction: if the interface is assumed to be fully ionized, cosmic-ray streaming speeds will be very low, forcing a buildup in pressure, while changing ionization would induce ionization-dependent transport and a drop in cosmic-ray energy at the same temperature band. As we go to lower temperatures, the $f_{\text{ion}}^{\text{min}} = 1.0$ curves continue to drop as slow-moving cosmic rays are consumed by collisions, but the $f_{\text{ion}}^{\text{min}} = 10^{-4}$ curves drop more gradually, as cosmic rays uniformly fill cold cloud interiors.

Lowering the magnetic field strength gives the most dramatic change. Cosmic rays streaming along the now-warped magnetic field lines preferentially squeeze through the underdense channels between clouds. The cosmic-ray and gamma-ray concentrations are both suppressed at low temperatures. Cosmic rays also do not occupy the high-temperature space: this low-density gas gets pushed into the halo, forced out by the




Figure 15. Accumulated collisional and collisionless cosmic-ray energy loss within one scale height (250 pc) for each of the mock ISM simulations. Different symbols represent changes in the transport model, while the color represents the ISM setup varying L , α , and B . The black dashed lines, both in the large figure and smaller zoom-in, show contours of constant total energy loss, and the solid black line shows the calorimetric limit, i.e., the total cosmic-ray energy injected into the box. Despite some significant differences in collisionless energy loss, especially when an additional, constant diffusion term is included, the collisional energy loss (and hence, diffuse gamma-ray emission) only varies within a factor of a few for most models. Changes in ISM setup and 2D versus 3D impart only small variations in total energy loss that are explained in the text.

cosmic-ray front that sweeps between clouds. In these low magnetic field cases, then, the great majority of cosmic rays lie within a narrow temperature range near the mean.

The bottom row of Figure 14 shows the PDF of hadronic gamma-ray emission $\propto e_{\text{CR}} n_{\text{gas}}$. The trends from the top row are clearly imprinted. For both transport models, gamma-ray emission is very concentrated near the mean density when $B = 1 \mu\text{G}$, but with $B = 5 \mu\text{G}$, a substantial fraction of the gamma-ray emission comes from cold clouds at temperatures $T < 10^4$ K. The low-temperature bins that dominate emission depend on the transport model. At intermediate densities $n \approx 1\text{--}10 \text{ cm}^{-3}$ (temperatures around 8×10^3 K), collisional energy losses are decreased considerably when ionization-dependent transport is included, with more energy loss occurring at the higher densities that fast-moving cosmic rays can leak into. In the fully ionized case, the highest densities are inaccessible to cosmic rays, as they lose energy instead in cloud interface layers ($T \approx 8 \times 10^3$ K) before penetrating into cloud cores. As we will see, this trade-off still leads to very similar total collisional loss rates (Section 6.4).

The spatial differences in collisional and collisionless energy loss are apparent in the bottom two panels of Figure 11. In both columns, most of the energy losses occur near the cosmic-ray source where the mean density is the highest. Compared to the fully ionized case, collisional losses are much more prevalent in dense gas, and collisionless energy loss not only occurs at the front cloud edge but also the back due to a second drop in v_A^{ion} that induces another steep cosmic-ray pressure gradient.

6.4. Total Cosmic-Ray Calorimetry

Figure 15 compiles the total collisional and collisionless energy loss from each of the mock ISM simulations we ran. Each symbol represents a different transport model, including a few simulations where we included an additional, constant cosmic-ray diffusion coefficient; this boosts cosmic-ray escape most drastically in the diffuse medium and cloud interfaces where the residence times are otherwise very long. Each color is a different ISM configuration, varying L , α , and B . Black dashed lines show contours of constant total cosmic-ray energy loss.

What one can see most clearly is that, although collisionless energy loss can vary by orders of magnitude, collisional energy loss only varies within a factor of a few regardless of transport model. All models without additional diffusion are clustered within the outlined box, meaning that, although the collisions occur in very different places (see Figure 14), the total collisional energy loss is about the same.

Zooming in on the clustered region of Figure 15, we do see some variation depending on ISM configuration and transport model, but the changes are within a factor of a few. For instance, while a lower magnetic field strength on average leads to lower transport speeds and longer cosmic-ray residence times in the ISM, the net collisional energy loss is comparable to that with a higher magnetic field strength. With weak magnetic fields, the collisionless energy loss $\propto v_A^{\text{ion}} \cdot \nabla P_{\text{CR}}$ decreases because $v_A^{\text{ion}} \propto B$ also decreases; the decrease is only by a factor of ≈ 2 , however, because slower transport speeds lead to a buildup in cosmic-ray pressure and higher ∇P_{CR} .

Varying α , which determines the density contrast, also makes a difference. With higher (lower) density contrast, the collisional energy loss $\propto n_{\text{gas}}$ increases (decreases) accordingly. Varying L has almost no consequence for collisional loss, however. The slight increase in collisions for $L=2$ compared to $L=5$ can be attributed to the slightly larger density range for $L=2$ (see Figure 10), which follows the same explanation as varying α .

We also ran three of our clump simulations in 3D, two of which had $f_{\text{ion}}^{\text{min}} = 1.0$ and one with $f_{\text{ion}}^{\text{min}} = 10^{-4}$. Each of these simulations had $B = 5 \mu\text{G}$, which, as we saw in the single-cloud simulations of Section 5.2, shows only small differences between 2D and 3D. That is reflected in Figure 15, as well, where the 2D and 3D energy losses are very similar. For larger cosmic-ray fluxes or larger plasma β , the differences between 2D and 3D may become more important, which we will explore in future work.

It is illuminating to also disentangle the differences in collisionless energy loss. If one were to draw a line diagonally across the plot from the lower left to upper right, that would show equal collisional and collisionless energy loss; we see that in all simulations run, collisional losses exceed collisionless losses, but this gap narrows when the density contrast (controlled by α) decreases and when ionization-dependent transport is included. Comparing, e.g., $L=2$, $\alpha=1.5$ simulations with different $f_{\text{ion}}^{\text{min}}$, we see that ionization-dependent transport decreases collisional energy loss but boosts collisionless energy loss, ending up with a similar total energy loss. The increase in collisionless energy loss can be explained by the second bottleneck that occurs at each back cloud edge (see e.g., Figure 4), therefore giving two interfaces that induce energy loss instead of just one. So ionization-dependent transport not only fails to decrease collisional losses substantially but can also offset this reduction by increasing collisionless energy loss—the net result is a very similar cosmic-ray calorimetry.

We find larger differences in collisionless and collisional energy loss when we include an additional diffusion term. Under the assumption of full ionization, an additional diffusion coefficient of $3 \times 10^{28} \text{ cm}^2 \text{ s}^{-1}$ significantly flattens the cosmic-ray pressure gradient and leads to a decrease in collisionless energy loss by a factor of tens. Interestingly, with $f_{\text{ion}}^{\text{min}} = 10^{-4}$, the same drop in energy loss can be achieved with a smaller diffusion coefficient of $3 \times 10^{27} \text{ cm}^2 \text{ s}^{-1}$. Without this additional diffusion, we saw that cosmic rays still

bottleneck at the cloud edges, which severely decreases the diffusive flux through the clouds. This bottleneck is effectively wiped out when even a small constant diffusivity is added, as it smooths the cosmic-ray pressure gradient and prevents the generation of confining waves; this maintains a large diffusivity through the partially neutral cloud interior, and transport no longer simply follows the ion Alfvén speed. The outcome is that cosmic rays do not sense the cloud and effectively free-stream through it. While this localized faster transport decreases the collisionless energy loss by orders of magnitude, the decrease in collisional loss is only a factor of a few because collisions primarily come from cosmic rays with long residence times in the warm ionized gas (see, e.g., Figure 14).

This finding is consistent with previous arguments in Chan et al. (2019) and Hopkins et al. (2021b), who compared full-galaxy simulations with varying diffusive and streaming transport with gamma-ray data. They found that, especially for dwarf galaxies, diffuse gamma-ray emission is far over-produced (by a factor of 10 or more) compared to observations of Local Group dwarf galaxies unless the cosmic-ray escape time is significantly reduced; despite strong outflows generated in these dwarf galaxies that would naively lead to fast advective escape, an additional cosmic-ray diffusion coefficient of $3 \times 10^{29} \text{ cm}^2 \text{ s}^{-1}$ is needed to lower emission to reasonable levels.

More specifically, Hopkins et al. (2021b) found that a significant fraction of gamma-ray emission was coming from the diffuse medium, where long cosmic-ray residence times compensated for the low gas density to enhance collisional losses $\propto e_{\text{CR}} n_{\text{gas}}$. The simulations that best fit observed gamma-ray and cosmic-ray grammage data, then, had to allow for fast cosmic-ray transport in the diffuse medium, while fast transport in cold, dense gas (where one would naively think most collisions occur) was not sufficient. Our study seems to corroborate this result, as we find collisional energy loss (hence, gamma-ray luminosity) to be fairly degenerate as we vary ISM setups and transport models. The largest decrease in collisions comes when an extra diffusion term is present—an additional simulation with $\kappa_{\parallel} = 3 \times 10^{29} \text{ cm}^2 \text{ s}^{-1}$ (not shown in Figure 15) continues the trend toward lower collisional energy loss.

7. Discussion

7.1. Comparison to Previous Bottleneck Simulations

While we have only simulated cosmic-ray to thermal pressure ratios of at most a few, therefore accelerating clouds to only a few km s^{-1} , we can already gleam some of the implications for galactic wind driving and compare to the broader literature on cloud “wind-tunnel” simulations and cosmic-ray bottlenecks. First, we have extended the work of Wiener et al. (2017a, 2019) with a faster, more accurate method of cosmic-ray transport (Jiang & Oh 2018) in the Athena++ code. This speed-up allowed us to, among other things, further explore the parameter space of magnetic field strengths and elucidate the implications of field line warping for cloud acceleration.

Our simulations without ionization-dependent transport are also similar to another recent simulation suite (Brüggen & Scannapieco 2020) but with some differences. Brüggen & Scannapieco (2020) simulate a cosmic-ray front with a ratio of the cosmic-ray pressure to thermal pressure ranging from a few

to a few tens, impinging upon a warm, fully ionized halo cloud with density $10^{-26} \text{ g cm}^{-3}$ and radius 100 pc, much larger and of lower density than our representative ISM cloud. Unlike our simulations, which define the cosmic-ray flux at the boundary for a finite time, they maintain a constant cosmic-ray energy density at the boundary for the entirety of their simulation. While we include cosmic-ray heating and collisional losses (which would be negligible for a diffuse halo cloud), they do not include heating but do include radiative cooling.

These differences in implementation and setup make it hard to draw fair comparisons. One differing result, however, is that we find a clear dependence of cloud acceleration on magnetic field strength, while they find only a small dependence (see, e.g., Figure 8 of Brüggen & Scannapieco 2020). This is likely because (1) they hold the boundary cosmic-ray energy density fixed, while our upstream cosmic-ray pressure depends on the Alfvén speed, and (2) they implement a switch that constrains cosmic rays to accelerate only the cells with density greater than 10% above the ambient density. If implemented in our simulations, these differences would change our results, as the amplitude and duration of the upstream cosmic-ray pressure gradient would have no magnetic field dependence, and the acoustic pulse that precedes the cosmic-ray front would not be present. As in Wiener et al. (2017a, 2019), we find that when $c_s > v_A$, this acoustic pulse plays an important role in “prepping” the cloud for the cosmic-ray front: it begins to accelerate the cloud, sometimes pushing it tens of parsecs before the cosmic-ray front reaches, and it affects the path of streaming cosmic rays by warping the magnetic field lines around the cloud. In our simulations with multiple clumps, especially those with low magnetic field strength ($c_s \gg v_A$), these effects are evidently consequential, as the cosmic-ray front squeezes between the clouds and almost entirely evacuates the hot, intercloud gas from the disk (see Figure 12). The role and evolution of this ambient gas are clearly interesting and should be explored in future work.

The main purpose of our paper is to consider the additional transport effects in denser, partially neutral clouds that are most appropriate for the ISM but may also be applicable to molecular gas expelled into galaxy halos. Everett & Zweibel (2011) found that, because the cosmic-ray flux quickly becomes diffusive and flattens the cosmic-ray pressure gradient (both upstream and inside the cloud), the cosmic rays cannot exert pressure forces and accelerate the cloud. Instead, on the upstream of the cloud, they can only exert a pressure difference (similar to ram pressure).

This picture changes if ion–neutral damping is not very strong, which, in our simulations, occurs in the cloud interfaces. In those regions, cosmic-ray pressure gradients impart momentum to the cloud with an amplitude and duration dependent on the magnetic field strength. As shown in Wiener et al. (2017a, 2019) and Brüggen & Scannapieco (2020), this bottleneck can lead to cloud acceleration possibly to hundreds of km/s in galaxy halos if the cloud can stay intact, which may be reinforced by radiative cooling.

Including ionization-dependent transport, the cloud acceleration we find lies somewhere between these fully ionized bottleneck simulations and the expectations from Everett & Zweibel (2011), with cosmic rays still able to exert forces at cloud interfaces, rather than cloud interiors. The net differences in cloud momenta considering fully ionized transport versus ionization-dependent transport in partially neutral gas differ by

less than a factor of 2. While we have only explored cosmic-ray to thermal pressure ratios of order unity, we expect cosmic-ray confinement to be even stronger at higher ratios, where cosmic-ray pressure gradients can more effectively excite confining waves. These results suggest that cosmic rays, especially near cosmic-ray sources where the pressure gradient is large, can still accelerate multiphase gas directly out of the ISM, despite not having a pressure gradient in the cloud interior. This deserves future attention and a more galactic wind-focused study than the one we have presented here.

7.2. Limitations

This work focuses on the interplay between cosmic rays and cold clouds, taking a critical look at the roles of bottlenecks induced upstream and elevated transport speeds in partially neutral clouds. We focus on the self-confinement model of cosmic-ray streaming, with advective flux modified by varying ion Alfvén speed and diffusive flux resulting solely from ion–neutral damping. We have not, until this point in the paper, explored the role of other damping mechanisms or instabilities that suppress or enhance confinement. While we mainly leave this to future work, we did rerun a subset of our 1D cloud simulations including nonlinear Landau damping (NLLD), which results from thermal particles taking energy from interacting Alfvén waves. Following Hopkins et al. (2021b), the damping rate can be written as

$$\Gamma_{\text{NLLD}} = \left[\frac{(\gamma_{\text{CR}} - 1)\pi^{1/2}}{8} \left(\frac{c_s v_A^{\text{ion}}}{r_L l_{\text{CR}}} \right) \left(\frac{e_{\text{CR}}}{e_B} \right) \right]^{1/2}. \quad (15)$$

This is proportional to the cosmic-ray pressure gradient, suggesting that waves may be efficiently damped in the cloud interface region where a steep pressure gradient forms, but the diffusive flux is inversely proportional: $\kappa_{\parallel} \propto (\nabla P_{\text{CR}})^{-1/2}$ (Loewenstein et al. 1991). Because of this, our simulations show only a small additional diffusion coefficient $< 10^{25} \text{ cm}^2/\text{s}$ in the cloud interface, and the resulting cosmic-ray and cloud evolution is almost exactly the same whether we account for NLLD or not.

Another limitation, which we plan to explore in future work, is that these simulations do not include cooling or conduction. These play important roles in cloud survival during acceleration (Brüggen & Scannapieco 2016; Gronke & Oh 2018), and they also set the thermodynamic state of the gas and the width of intermediate temperature transition regions, or cloud envelopes. Note that a typical Fields length in 10^4 K gas is less than 0.1 pc, much smaller than our fiducial cloud setup with a 5 pc interface. Heating from cosmic-ray streaming, while it noticeably heats the interface in our simulations, is not likely to set the interface width in real, cooling clouds. From our simulations, we find collisionless loss rates (hence, heating rates) less than $10^{-25} \text{ erg cm}^{-3} \text{ s}^{-1}$, subdominant compared to the expected radiative cooling and conductive heating rates if those were taken into account (see also Figure 9 of Everett & Zweibel 2011).

Models and observations of molecular clouds exposed to interstellar radiation fields, however, do show a gradual transition between hot and cold phases, with warm, ionized envelopes surrounding the cold cores (Goldsmith & Li 2005). Indeed, a growing literature on cosmic-ray penetration into molecular clouds has come to similar conclusions about the

trapping of cosmic rays by either extrinsic or self-generated turbulence in interface regions (Morlino & Gabici 2015; Schlickeiser et al. 2016; Dogiel et al. 2018; Ivlev et al. 2018; Inoue 2019), especially near cosmic-ray sources where it appears that cosmic rays can be confined even in the presence of neutral particles (Nava et al. 2016; Brahimi et al. 2020).

7.3. Further Observational Constraints

Constraining these theoretical results with observations is important and already underway to truly diagnose cosmic-ray transport in dense, cold clouds. Silsbee & Ivlev (2019) considered whether the low-energy (less than GeV) cosmic-ray spectrum from the Voyager probe would be better fit by diffusive or free-streaming propagation of cosmic rays penetrating through a large column density into the Local Bubble. While confinement will be stronger at these energies than for GeV cosmic rays, they find that diffusive propagation is a better fit to the spectrum than free-streaming propagation, which would additionally require an unreasonably high column density to fit the Voyager data.

Fujita et al. (2021) model free-streaming and diffusive propagation (due to preexisting MHD turbulence) in molecular clouds, obtaining profiles of the ionization rate, 6.4 keV line flux, and gamma-ray emission. While there are only a few observed supernova remnants for which all three of these diagnostics are available, the qualitative comparison in their Discussion section finds that both diffusive and free-streaming propagation are likely realized.

Joubaud et al. (2020) studied a diffuse cloud near the Orion-Eridanus superbubble that was observed to have a 34% lower cosmic-ray flux compared to the local flux estimate. As the authors note, the estimates of Everett & Zweibel (2011) suggest only a slight drop in cosmic-ray flux within cold clouds and would require a much higher than observed magnetic field strength to account for such a 30% drop. Our simulations assume cloud parameters closer to the Eridu estimates ($n \approx 7 \text{ cm}^{-3}$ and $B \approx 5 \text{ } \mu\text{G}$) and find that the cosmic-ray pressure drops by a factor of ≈ 2 if the cloud is partially ionized and steadily drops much more if the cloud is fully ionized. If the cloud lies along a magnetic flux tube with cosmic rays streaming into the halo, as speculated by Joubaud et al. (2020), then the cloud may even be accelerating due to a cosmic-ray bottleneck. In any case, if a warm ionized envelope exists, in which the Alfvén speed drops and reinforces the cosmic-ray pressure gradient, it is not unreasonable, given our results, for the cloud-interior cosmic-ray pressure to be 30% lower than in the surrounding medium. More observational studies in this vein will be very helpful in further constraining cosmic-ray-cloud interactions.

8. Conclusions

Many MHD + cosmic-ray simulations, aimed at modeling galaxy evolution, now take into account the self-confinement of GeV-energy cosmic rays through frequent scattering off cosmic-ray-generated Alfvén waves. While this kinetic process occurs at scales of order the cosmic-ray gyroradius, many orders of magnitude below the current resolution limit of galaxy evolution simulations, fluid theories of this streaming transport have been developed and, with the advent of various numerical techniques, implemented in MHD codes for use in a galaxy-scale context. With these tools now available, our

understanding of cosmic-ray influence in the interstellar, circumgalactic, and intracluster medium has rapidly progressed; however, as of this writing, two important trends are evident in most published simulations: (1) it is assumed that GeV-energy cosmic-ray transport proceeds as if the background thermal gas is everywhere fully ionized. In reality, the waves the cosmic rays resonate with propagate only in the ions, and streaming is therefore at the ion Alfvén speed $v_A^{\text{ion}} = v_A / \sqrt{f_{\text{ion}}}$. Also, ion-neutral friction can damp the confining waves, thereby decoupling cosmic rays and greatly increasing their diffusive flux. (2) Partially neutral ISM clouds and other density irregularities that induce decoupling and nonlinear cosmic transport are frequently underresolved. If we imagine the multiphase ISM as an obstacle course of coupled versus decoupled regimes, further complicated by hadronic and Coulomb collisions between cosmic rays and ambient gas, then how do cosmic rays navigate this medium, and can we constrain this transport with observations? The answers have implications for ISM dynamics, interpretation of observations, and the acceleration of cold interstellar gas in galactic winds and fountains.

In this paper, we presented a suite of high-resolution, idealized simulations of cosmic-ray fronts navigating the multiphase ISM. We ran each simulation twice: first, assuming the medium is fully ionized and, second, accounting for changes in ionization that induce “ionization-dependent transport”. We began with 1D simulations of an energetically significant cosmic-ray front hitting a single, partially neutral cloud with radius and density typical of ISM conditions. We then explored how cosmic-ray transport, cosmic-ray energy loss, and cloud acceleration and morphology varied in multiple dimensions and varying magnetic field strength. We then ran the same cosmic-ray fronts through mock ISM setups: slabs of lognormally distributed multiphase clumps on top of an exponentially decaying mean density. We assessed how cosmic rays sampled the ISM and how ionization-dependent transport affected cloud acceleration and collisional and collisionless energy loss.

Our main findings were:

1. Cloud interfaces can play a crucial role in setting the cosmic-ray propagation through the cloud interior. The drop in Alfvén speed while the gas is still fully ionized leads to a bottleneck in the interface; this reinforces the cosmic-ray pressure gradient that generates confining waves and locks cosmic rays to the wave frame. The effects of ion-neutral damping are thereby suppressed, and the advective flux rather than diffusive flux generally dominates transport ($v_{\text{st}} \approx v_A^{\text{ion}}$).
2. The staircase structure that develops at cloud interfaces facilitates collisionless energy transfer from the cosmic rays to waves $\propto v_A^{\text{ion}} \nabla P_{\text{CR}}$. Unlike for a fully ionized cloud where a single cosmic-ray gradient forms (Wiener et al. 2017a, 2019), there is now a bottleneck at both the front and back cloud edges. In our mock ISM simulations, this extra energy loss partially offsets a reduction in collisional energy loss, keeping the total cosmic-ray calorimetry similar to when full ionization is assumed. The additional gas heating may also have implications for ion abundances and kinematics at cloud interfaces, though we estimate that the role of heating would be secondary to cooling and conduction. We plan to explore this further by more self-consistently modeling interface

properties rather than leaving the interface width as a free parameter.

3. Cosmic rays sample the ISM differently when varying ionization is accounted for. Ionization-dependent transport allows cosmic rays to leak into dense cloud interiors while cosmic rays are confined to thin cloud boundary layers, assuming full ionization. These differences are imprinted on gamma-ray emission maps.
4. Total collisional losses are very similar for the two transport models. With ionization-dependent transport, short cosmic-ray residence times in clouds are offset by the denser gas they access. So while gamma-ray emission maps with the two transport models look very different, the gamma-ray luminosity is dominantly set by the cosmic-ray transport in the warm ionized medium, which is the same in each transport model. This result is consistent with Hopkins et al. (2021b).
5. When ionization is accounted for, cloud acceleration is decreased because ionization-dependent transport in cloud interiors flattens the cosmic-ray pressure profile. The change is not as large as one might expect, however, as our simulations exhibit a factor of <2 difference in cold gas momentum. Cosmic rays still bottleneck at the cloud edges, leading to pressure gradient forces that push the cloud. This suggests that, contrary to previous expectation, ionization-dependent transport does not greatly inhibit cosmic rays from accelerating cold, molecular gas in galactic outflows. The magnetic field topology, which can be warped upstream by strong cosmic-ray fronts, may be a greater inhibitor of direct cosmic-ray acceleration of clouds. We will explore these aspects more thoroughly in a future study.

Clearly, cosmic-ray propagation through the multiphase ISM is a very rich problem with fundamental implications for ISM dynamics and feedback. To understand how cosmic rays shape their environments, we need to continue applying first-principles cosmic-ray transport models to astrophysical problems, test the outcomes against observations, and use those constraints to create more complete models of cosmic-ray transport. Simulations that employ streaming at the gas Alfvén speed will drastically underestimate transport speeds in partially neutral clouds, thereby overestimating cloud acceleration. On the other hand, constant field-aligned diffusion models have no dependence on local gas conditions, thereby precluding bottleneck formation at cloud interfaces, and they omit the energy transfer from cosmic rays to thermal gas, fundamentally changing cosmic-ray influence on galaxies (Wiener et al. 2017b). “Two- κ ” prescriptions (e.g., Farber et al. 2018), while tying the transport speed to local gas temperature, again prohibit bottlenecks at cloud interfaces and may overestimate the effective transport speed depending on the choice of κ . These models, then, likely underestimate the dynamical influence of cosmic rays on partially neutral gas.

In light of our results, we instead advocate that galaxy evolution simulations should, to zeroth order, model streaming at the ion Alfvén speed with a temperature-based ion fraction. This method will more fully capture the generation of bottlenecks and steep pressure gradients at cloud interfaces, as well as ionization-dependent transport in cloud interiors that is robust, at least in these simulations, to changes in ionization fraction. More detailed modeling of ionization from local

sources and lower energy cosmic rays could alter this picture and should be a next step.

The authors are greatly indebted to Yan-Fei Jiang for providing us with his two-moment method of cosmic-ray transport in Athena++. We also thank the anonymous referee whose comments greatly improved the quality and clarity of this manuscript, as well as Josh Wiener, Mateusz Ruszkowski, Ryan Farber, Peng Oh, and Max Gronke for stimulating discussions. This research was supported in part by the National Science Foundation under grant No. NSF PHY-1748958 and by the Gordon and Betty Moore Foundation through grant No. GBMF7392. EGZ is happy to acknowledge support from NSF Grant AST 2007323. This work made extensive use of the *yt* (Turk et al. 2011) software package, a toolkit for analyzing and visualizing quantitative data. Computations were performed using the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562 (Towns et al. 2014). Specifically, our computing resources stemmed from allocation TG-AST190019 on the Stampede2 supercomputer.

Appendix A Effect of Resolution

In this section, we carry out a small resolution study, primarily in 1D, to test convergence and inform larger-scale simulations of galaxy evolution, which may underresolve cosmic-ray interactions with multiphase structures. Figure 16 shows our fiducial 1D cloud of radius 10 pc and interface width 5 pc simulated at resolutions of 2, 1, 0.5, and 0.25 pc. Lower resolution leads to a smaller total drop in cosmic-ray energy through the cloud, less heating at the front edge, and different cloud morphologies (note the disappearance of the upstream “thumb” caused by heating). Convergence is achieved at a resolution of 0.5 pc, which resolves the 5 pc-thick interface by a sufficient number of cells. Similarly, our simulations with a 10 pc interface width reached convergence at 1 pc.

Because the total pressure drop determines the force on the cloud, we expect that underresolved simulations will underestimate cloud acceleration. To test this, we ran our 2D cloud simulations (with $f_{\text{ion}}^{\text{min}} = 10^{-4}$) at resolutions of 4, 2, and 1 pc. The resulting cloud momenta and velocities are shown in Figure 17. As expected, deteriorating resolution decreases cloud acceleration and, hence, the perceived effectiveness of cosmic-ray-driven feedback.

While this parsec or subparsec resolution requirement is initially daunting, it is worth noting that, if transport is dominated by the advective rather than diffusive flux, such problems become more computationally tractable. Our choice to set a high $V_m = 10^{10} \text{ cm s}^{-1}$, advective cap of 10^9 cm s^{-1} , and diffusive cap of $3 \times 10^{30} \text{ cm}^2 \text{ s}^{-1}$ in the first place was a preemptive measure anticipating that large diffusive fluxes, in particular, would need to be accommodated. As we found, that is rarely the case, and transport speeds are generally limited by v_A^{ion} . We found only small differences, then, when V_m and the diffusive and advective caps are each decreased by a factor of 5–10, which should provide a glimmer of hope for convergence of future large-scale simulations with ionization-dependent transport.

Finally, we test how cosmic rays sample the ISM and transfer energy in our clumpy simulations with lower




Figure 16. Density and cosmic-ray energy density for the fiducial cosmic-ray front impinging on a cloud of radius 10 pc and interface width 5 pc. Each panel corresponds to a different resolution. When the interface region is well resolved (≈ 10 cells per interface), the results are well converged. With poorer resolution, the cosmic-ray front does not respond to the drop in Alfvén speed at the interface, therefore barreling through the cloud and appearing on the other side with a larger cosmic-ray energy. The spike in density on the front cloud edge, caused by cosmic-ray heating, is also not as noticeable at low resolution.




Figure 17. Cloud momentum (left) and velocity (right) at varying resolutions of 1 pc (solid lines), 2 pc (dashed lines), and 4 pc (dotted-dashed lines) for the 2D cloud simulation of Section 5.2 with $B = 5 \mu\text{G}$ (black lines) and $B = 1 \mu\text{G}$ (red lines). Deteriorating resolution leads to less cloud acceleration because cosmic-ray pressure gradients are smoothed at the unresolved cloud boundary. Even at the fiducial 1 pc resolution, we do not appear to be reaching convergence, as expected from Figure 16.




Figure 18. Probability distribution function (see also Figure 14) of cosmic-ray energy density (left) and histogram of gamma-ray luminosity (right) at varying resolutions of 2 pc (solid lines), 4 pc (dashed lines), 8 pc (dotted-dashed lines), and 32 pc (dotted lines). For either magnetic field strength tested, lowering resolution leads to higher cosmic-ray occupancy in cold gas, especially at interface temperatures $T \approx 10^4$ K. This leads to slightly more gamma-ray emission coming from cold clumps. This is especially noticeable in the $B = 1 \mu\text{G}$ simulations, though the majority of emission still comes from the $T \approx 2 \times 10^4$ K band. Note the change in x-axis limits in the right panel to focus on the lower temperature region where most gamma-ray emission comes from.

resolutions. We again focus on ionization-dependent transport with $f_{\text{ion}}^{\text{min}} = 10^{-4}$, and we test resolutions of 32, 8, 4, and 2 pc in our $(L, \alpha) = (2, 1.5)$ setups. The resulting cosmic-ray energy density PDFs and gamma-ray luminosity histograms are shown

in Figure 18. Lowering the resolution smooths cloud interfaces, allowing cosmic rays to leak more easily into cold, dense gas. This trend is evident in both the $B = 5 \mu\text{G}$ and $B = 1 \mu\text{G}$ simulations. This leads to more gamma-ray emission coming

from the dense, cold gas, especially in the interface temperature range near 8×10^3 K.

The total gamma-ray luminosities over all temperature bins are relatively unchanged when the resolution is varied, though, and even appear to decrease in the $B = 5 \mu\text{G}$ case with 32 pc resolution. This is consistent, also, with our findings of Section 6 where we changed the smoothness of the clumpy ISM via the α parameter. In Figure 15, one can see that lower α , akin to lower resolution that would smooth out density perturbations, actually leads to a decrease in collisional losses. Poor resolution, then, appears not to be the dominant factor that grossly overestimates gamma-ray emission in large-scale simulations, but more high-resolution simulations of cosmic rays in different environments (e.g., superbubbles versus mean ISM) and for a range of mean gas densities are needed.

Appendix B Validity of a Fluid Cosmic-Ray Model

A possible limitation of this work concerns the validity of a fluid model for cosmic rays. This assumption rests upon the cosmic-ray mean free path λ_{mfp} being shorter than other length scales of interest, namely the cosmic-ray scale length l_{CR} , and is, in fact, baked into the diffusivity we have implemented: Assuming a short mean free path and balancing scattering against acceleration down the pressure gradient leads to a scattering rate $\nu \sim c^2/(l_{CR}v_s)$, where v_s is the streaming velocity at which wave damping balances wave growth. This implies that the ratio of mean free path to scale length $\lambda_{mfp}/l_{CR} \sim v_s/c < 1$. So even at cloud interfaces where the diffusivity is initially very large and the scale length becomes




Figure 19. The cosmic-ray mean free path $\lambda_{mfp} = \kappa_{||}/c$ versus the cosmic-ray scale length l_{CR} for the 1D simulations of Section 5.1.

very short as the cosmic-ray front approaches, $\lambda_{mfp} < l_{CR}$ by design.

In Figure 19, we plot the mean free path $\lambda_{mfp} = \kappa/c$ versus $l_{cr} = P_{CR}/\nabla P_{CR}$ at a time of 28 Myr when the cosmic-ray front is just hitting the cloud we presented in Section 5.1. At every snapshot we output to data, the mean free path is below the cosmic-ray scale length, but clearly, as in this snapshot, there are locations upstream of the cosmic-ray front where the mean free path is greater than the scale length a few cells away from it.

Appendix C

Cosmic-Ray Exclusion via Collisional Losses

One of the main findings of our work is that the interface between hot and cold gas plays a crucial role in cosmic-ray propagation and dynamical effects on the cloud. Our result is broadly consistent with a growing literature on cosmic-ray penetration into molecular clouds, the history of which is quite rich (see, e.g., Skilling & Strong 1976 and Cesarsky & Volk 1978 for seminal papers) and outlined in Everett & Zweibel (2011). Specifically, recent works have noted that cosmic-ray penetration into molecular clouds is highly non-linear and self-modulated by the excitation of waves in the cloud interface (see the Discussion section). While we have focused on a cosmic-ray front moving with a preferred direction toward a cloud, it may be that molecular clouds sit in a “sea” of cosmic rays—an initially flat cosmic-ray energy profile. In this case, with no preexisting pressure gradient, the cosmic-ray profile is modulated by collisional loss rates in the dense molecular cloud. The resulting pressure difference sucks fresh cosmic rays into the cloud, and how much those exterior cosmic rays can pressurize the cloud interior depends on the transport through the cloud interface.

We model this scenario for two different clouds each of radius 50 pc, with varying interface widths of 25, 5, and 0.5 pc, and with densities of 10 and 100 cm^{-3} . We use a uniform resolution of 0.5 pc for simulations with interface widths 25 and 5 pc, and we increase the resolution to 0.05 pc for the 0.5 pc interface run in order to keep the interface resolved. Initially, the cosmic-ray energy density is uniform everywhere, but it drops within the cloud as collisions decrease the cosmic-ray pressure. We use inflow/outflow boundaries to allow for a steady influx of cosmic rays into the simulation and back into the cloud.

Figure 20 shows our results at snapshots when the cosmic-ray energy density has “bottomed out”, i.e., when a steady-state has been reached between cosmic-ray inflow through the interface and energy loss in the interior. The cloud morphology changes somewhat due to the pressure imbalance, but the cosmic-ray pressure we start with is at least a few times smaller than the thermal pressure, so the effect is small. For the fiducial cloud density of 10 cm^{-3} (top row) with minimum ion fractions of 10^{-4} and 10^{-1} , the cosmic-ray energy stays fairly level. As expected, the interface is the biggest bottleneck for inflowing cosmic rays. In the $f_{\text{ion}}^{\text{min}} = 10^{-4}$ simulations, a thicker interface leads to a larger drop in cosmic-ray energy within the cloud. For $f_{\text{ion}}^{\text{min}} = 10^{-1}$, the diffusive flux term becomes more important,




Figure 20. Profiles of 1D clouds with peak densities of 10 cm^{-3} (top row) and 100 cm^{-3} (bottom row) initially bathed in a sea of cosmic rays with uniform pressure. Collisional losses decrease cosmic-ray pressure within the cloud core, and the steady-state pressure in the cloud is governed by a balance between this loss rate and the influx of new cosmic rays through the interface, which we vary in width from 25 pc (solid lines), 5 pc (dashed lines), and 0.5 pc (dotted-dashed lines). For the 10 cm^{-3} cloud, the cosmic-ray pressure within the cloud is almost the same as the outside pressure, especially for the $f_{\text{ion}}^{\text{min}} = 10^{-4}$ simulations with interfaces too thin to instigate significant bottlenecks. Wider and denser clouds show greater drops in cosmic-ray pressure.

and the trend with interface width becomes more complicated, though the differences in cosmic-ray pressure are quite small.

For a denser cloud (100 cm^{-3}), which is the fiducial value in Everett & Zweibel (2011), the energy drop is more substantial. There is a slight trend toward higher cosmic-ray energy inside the cloud when the interface is thinner, but generally, the energy drops are very similar for different interface widths and different initial cosmic-ray energies as long as the interface is well resolved. It is worth noting that, while we try to choose steady-state snapshots, the cloud interface is constantly changing due to collisionless heating, which in these simulations is not counteracted by radiative cooling. In the high initial cosmic-ray energy cases, where the magnitude of this heating is greatest, the interface eventually widens and shows clear density ripples moving outwards from the cloud. This can throw the cosmic-ray energy out of steady-state, so we have chosen early (within tens of Myr) snapshots with seemingly steady-state profiles for these figures. Closer inspection of how the interface changes in the additional presence of radiative cooling and thermal conduction is left to future work.

This study addresses the assertion in Everett & Zweibel (2011) that cosmic-ray energy will be roughly uniform inside and outside cold clouds and shows that this very much depends on interface properties. It also shows that the extent to which cosmic rays penetrate clouds and the relative importance of diffusive and advective fluxes depends somewhat on the cosmic-ray content already in the cloud. This is different from the single cosmic-ray front simulations we have focused on in this paper, where the presence of a steep front very much negates the diffusive flux and locks cosmic rays to Alfvén waves. In the now-uniform pressure setup, the collisional loss time is not necessarily short enough to cause a steep pressure gradient, especially when the interface is narrow, and the diffusive flux can be significant. Future work will address more realistic scenarios with multiple cosmic-ray bursts from different directions, as well as cosmic-ray production within cold clouds.

ORCID iDs

Chad Bustard  <https://orcid.org/0000-0002-8366-2143>

Ellen G. Zweibel  <https://orcid.org/0000-0003-4821-713X>

References

- Ackermann, M., Albert, A., Atwood, W. B., et al. 2016, *A&A*, **586**, A71
 Ajello, M., Di Mauro, M., Paliya, V. S., & Garrappa, S. 2020, *ApJ*, **894**, 88
 Amato, E., & Blasi, P. 2018, *AdSpR*, **62**, 2731
 Banda-Barragán, W., Brüggen, M., Heesen, V., et al. 2020, arXiv:2011.05240
 Becker Tjus, J., & Merten, L. 2020, *PhR*, **872**, 1
 Blasi, P., & Amato, E. 2019, *PhRvL*, **122**, 051101
 Boettcher, E., Zweibel, E. G., Yoast-Hull, T. M., & Gallagher, J. S. 2013, *ApJ*, **779**, 12
 Boulares, A., & Cox, D. P. 1990, *ApJ*, **365**, 544
 Brahimi, L., Marcowith, A., & Ptuskin, V. S. 2020, *A&A*, **633**, A72
 Breitschwerdt, D., McKenzie, J. F., & Voelk, H. J. 1991, *A&A*, **245**, 79
 Brüggen, M. 2013, *MNRAS*, **436**, 294
 Brüggen, M., & Scannapieco, E. 2016, *ApJ*, **822**, 31
 Brüggen, M., & Scannapieco, E. 2020, *ApJ*, **905**, 19
 Buck, T., Pfrommer, C., Pakmor, R., Grand, R. J. J., & Springel, V. 2020, *MNRAS*, **497**, 1712
 Bustard, C., Zweibel, E. G., D’Onghia, E., Gallagher, J. S., & Farber, R. I. 2020, *ApJ*, **893**, 29
 Butsky, I. S., & Quinn, T. R. 2018, *ApJ*, **868**, 108
 Cesarsky, C. J., & Volk, H. J. 1978, *A&A*, **70**, 367
 Chan, T. K., Kereš, D., Hopkins, P. F., et al. 2019, *MNRAS*, **488**, 3716
 Crocker, R. M., Krumholz, M. R., & Thompson, T. A. 2021, *MNRAS*, **502**, 1312
 Dashyan, G., & Dubois, Y. 2020, *A&A*, **638**, A123
 De Pontieu, B., Martens, P. C. H., & Hudson, H. S. 2001, *ApJ*, **558**, 859
 Dogiel, V. A., Chernyshov, D. O., Ivlev, A. V., et al. 2018, *ApJ*, **868**, 114
 Enßlin, T. A., Pfrommer, C., Springel, V., & Jubelgas, M. 2007, *A&A*, **473**, 41
 Everett, J. E., & Zweibel, E. G. 2011, *ApJ*, **739**, 60
 Everett, J. E., Zweibel, E. G., Benjamin, R. A., et al. 2008, *ApJ*, **674**, 258
 Farber, R., Ruszkowski, M., Yang, H. Y. K., & Zweibel, E. G. 2018, *ApJ*, **856**, 112
 Fu, L., Xia, Z. Q., & Shen, Z. Q. 2017, *MNRAS*, **471**, 1737
 Fujita, Y., Nobukawa, K. K., & Sano, H. 2021, *ApJ*, **908**, 136
 Girichidis, P., Naab, T., Hanasz, M., & Walch, S. 2018, *MNRAS*, **479**, 3042
 Goldsmith, P. F., & Li, D. 2005, *ApJ*, **622**, 938
 Gronke, M., & Oh, S. P. 2018, *MNRAS*, **480**, L111
 Gronke, M., & Oh, S. P. 2020, *MNRAS*, **492**, 1970
 Hanasz, M., Lesch, H., Naab, T., et al. 2013, *ApJL*, **777**, L38
 Heintz, E., Bustard, C., & Zweibel, E. G. 2020, *ApJ*, **891**, 157
 Holguin, F., Ruszkowski, M., Lazarian, A., Farber, R., & Yang, H. Y. K. 2019, *MNRAS*, **490**, 1271
 Hollenbach, D. J., & Tielens, A. G. G. M. 1999, *RvMP*, **71**, 173
 Hopkins, P. F., Chan, T. K., Garrison-Kimmel, S., et al. 2020c, *MNRAS*, **492**, 3465
 Hopkins, P. F., Chan, T. K., Ji, S., et al. 2021a, *MNRAS*, **501**, 3640
 Hopkins, P. F., Squire, J., Chan, T. K., et al. 2021b, *MNRAS*, **501**, 4184
 Inoue, T. 2019, *ApJ*, **872**, 46
 Ipavich, F. M. 1975, *ApJ*, **196**, 107
 Ivlev, A. V., Dogiel, V. A., Chernyshov, D. O., et al. 2018, *ApJ*, **855**, 23
 Ji, S., Chan, T. K., Hummels, C. B., et al. 2020, *MNRAS*, **496**, 4221
 Jiang, Y.-F., & Oh, S. P. 2018, *ApJ*, **854**, 5
 Joubaud, T., Grenier, I. A., Casandjian, J. M., Tolsdorf, T., & Schlickeiser, R. 2020, *A&A*, **635**, A96
 Kanjilal, V., Dutta, A., & Sharma, P. 2021, *MNRAS*, **501**, 1143
 Kempster, P., & Quataert, E. 2020, *MNRAS*, **493**, 1801
 Kim, C.-G., Ostriker, E. C., Somerville, R. S., et al. 2020, *ApJ*, **900**, 61
 Klein, R. I., McKee, C. F., & Colella, P. 1994, *ApJ*, **420**, 213
 Krumholz, M. R., Crocker, R. M., Xu, S., et al. 2020, *MNRAS*, **493**, 2817
 Kulsrud, R., & Pearce, W. P. 1969, *ApJ*, **156**, 445
 Lacki, B. C., Thompson, T. A., Quataert, E., Loeb, A., & Waxman, E. 2011, *ApJ*, **734**, 107
 Lazarian, A. 2016, *ApJ*, **833**, 131
 Loewenstein, M., Zweibel, E. G., & Begelman, M. C. 1991, *ApJ*, **377**, 392
 Lopez, L. A., Auchettl, K., Linden, T., et al. 2018, *ApJ*, **867**, 44
 Mao, S. A., & Ostriker, E. C. 2018, *ApJ*, **854**, 89
 Morlino, G., & Gabici, S. 2015, *MNRAS*, **451**, L100
 Nava, L., Gabici, S., Marcowith, A., Morlino, G., & Ptuskin, V. S. 2016, *MNRAS*, **461**, 3552
 Neufeld, D. A., & Wolfire, M. G. 2017, *ApJ*, **845**, 163
 O’C Drury, L., Duffy, P., & Kirk, J. G. 1996, *A&A*, **309**, 1002
 Pfrommer, C., Pakmor, R., Schaal, K., Simpson, C. M., & Springel, V. 2017, *MNRAS*, **465**, 4500
 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., & Corlies, L. 2016, *MNRAS*, **456**, 582
 Schlickeiser, R., Caglar, M., & Lazarian, A. 2016, *ApJ*, **824**, 89
 Sharma, P., Colella, P., & Martin, D. F. 2009, arXiv:0909.5426
 Silsbee, K., & Ivlev, A. V. 2019, *ApJ*, **879**, 14
 Skilling, J. 1971, *ApJ*, **170**, 265
 Skilling, J., & Strong, A. W. 1976, *A&A*, **53**, 253
 Stone, J. M., Tomida, K., White, C. J., & Felker, K. G. 2020, *ApJS*, **249**, 4
 Sutherland, R. S., & Dopita, M. A. 1993, *ApJS*, **88**, 253
 Tang, Q.-W., Wang, X.-Y., & Tam, P.-H. T. 2014, *ApJ*, **794**, 26
 Thomas, T., & Pfrommer, C. 2019, *MNRAS*, **485**, 2977
 Towns, J., Cockerill, T., Dahan, M., et al. 2014, *CSE*, **16**, 62
 Turk, M. J., Smith, B. D., Oishi, J. S., et al. 2011, *ApJS*, **192**, 9
 Uhlig, M., Pfrommer, C., Sharma, M., et al. 2012, *MNRAS*, **423**, 2374
 Wang, X., & Fields, B. D. 2018, *MNRAS*, **474**, 4073
 Wentzel, D. G. 1968, *ApJ*, **152**, 987
 Wentzel, D. G. 1969, *ApJ*, **156**, 303
 Wiener, J., Oh, S. P., & Zweibel, E. G. 2017a, *MNRAS*, **467**, 646
 Wiener, J., Pfrommer, C., & Oh, S. P. 2017b, *MNRAS*, **467**, 906
 Wiener, J., Zweibel, E. G., & Ruszkowski, M. 2019, *MNRAS*, **489**, 205
 Yan, H., & Lazarian, A. 2004, *ApJ*, **614**, 757
 Yoast-Hull, T. M., Gallagher, J. S., & Zweibel, E. G. 2015, *MNRAS*, **453**, 222
 Yoast-Hull, T. M., Gallagher, J. S., & Zweibel, E. G. 2016, *MNRAS*, **457**, L29
 Zweibel, E. G. 2017, *PhPI*, **24**, 055402
 Zweibel, E. G. 2020, *ApJ*, **890**, 67