

Cosmic ray interstellar propagation tool using Itô Calculus (CRIPTIC): software for simultaneous calculation of cosmic ray transport and observational signatures

Mark R. Krumholz^{1,2}★, Roland M. Crocker¹ and Matt L. Sampson¹

¹Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2601, Australia

²ARC Centre of Excellence for Astronomy in Three Dimensions (ASTRO-3D), Canberra, ACT 2601, Australia

Accepted 2022 September 17. Received 2022 August 28; in original form 2022 July 27

ABSTRACT

We present CRIPTIC, the Cosmic Ray Interstellar Propagation Tool using Itô Calculus, a new open-source software package to simulate the propagation of cosmic rays through the interstellar medium and to calculate the resulting observable non-thermal emission. CRIPTIC solves the Fokker–Planck equation describing transport of cosmic rays on scales larger than that on which their pitch angles become approximately isotropic, and couples this to a rich and accurate treatment of the microphysical processes by which cosmic rays in the energy range $\sim$ MeV to $\sim$ PeV lose energy and produce emission. CRIPTIC is deliberately agnostic as to both the cosmic ray transport model and the state of the background plasma through which cosmic rays travel. It can solve problems where cosmic rays stream, diffuse, or perform arbitrary combinations of both, and the coefficients describing these transport processes can be arbitrary functions of the background plasma state, the properties of the cosmic rays themselves, and local integrals of the cosmic ray field itself (e.g. the local cosmic ray pressure or pressure gradient). The code is parallelized using a hybrid OpenMP-MPI paradigm, allowing rapid calculations exploiting multiple cores and nodes on modern supercomputers. Here, we describe the numerical methods used in the code, our treatment of the microphysical processes, and the set of code tests and validations we have performed.

Key words: radiation mechanisms: non-thermal – methods: numerical – cosmic rays.

1 INTRODUCTION

The last few years have seen an explosion of interest in cosmic rays (CRs) from two distinct angles. The star- and galaxy-formation communities have engaged in intense study of CRs as a form of stellar feedback in star and galaxy formation (e.g. Enßlin et al. 2007; Socrates, Davis & Ramirez-Ruiz 2008; Uhlig et al. 2012; Salem & Bryan 2014; Girichidis et al. 2016; Wiener, Pfrommer & Oh 2017; Chan et al. 2019; Hopkins et al. 2020; Crocker, Krumholz & Thompson 2021a, b), while the high-energy astrophysics community has paid increasing attention to star-forming galaxies and the CRs within them as important sources at both radio and γ -ray wavelengths (e.g. Thompson et al. 2006; Lacki & Thompson 2010; Lacki, Thompson & Quataert 2010; Yoast-Hull, Gallagher & Zweibel 2016; Peretti et al. 2019; Krumholz et al. 2020; Roth et al. 2021; Werhahn et al. 2021a; Werhahn, Pfrommer & Girichidis 2021b; Hopkins et al. 2022). However, studies in both of these areas are hampered by our lack of understanding of the fundamental plasma processes by which CRs couple to the background gas in galaxies through which they propagate. While most plasma physics models predict that CRs should be self-confined at relatively low energy and confined by extrinsic turbulence at high energy (e.g. Zweibel 2017, and references therein), the energy at which this transition occurs, and the normalization and energy dependence of the rate of transport in each of these two regimes, remains fundamentally

uncertain (and likely dependent on the environment), and simple models have proven challenging to reconcile with observational constraints (e.g. Hopkins et al. 2021a, b). Until we gain a better understanding of how CRs couple to the gas in galaxies, it will be difficult to make definitive statements about either the role of CRs in regulating galaxy formation or the contribution of star-forming galaxies to the non-thermal sky.

One promising avenue for progress in this area is making detailed comparisons between the predictions made by different CR transport models and observations of galaxies’ non-thermal emission. During their transport through gas, CRs suffer repeated collisions with the molecules, atoms, and nuclei they encounter. Low-energy CRs ionize the gas, altering its chemistry. High-energy CRs collide directly with gas nuclei, producing sprays of unstable secondary particles that decay almost immediately into final-state particles, including γ -ray photons, neutrinos, and relativistic electrons and positrons. The γ -rays and neutrinos are (in principle) directly observable, while the electrons and positrons go on to produce their own radiative signatures at radio and γ -ray wavelengths. These signatures hold open the possibility of distinguishing between CR transport models, because different models predict different behaviours as a function of CR energy and galactic environment, which in turn manifest as changes in galaxies’ non-thermal spectra (e.g. Krumholz et al. 2020; Crocker et al. 2021a; Ambrosone et al. 2022).

Our observational knowledge of these signatures will expand radically in the next few years as new instruments come online. The Cherenkov Telescope Array (CTA; Cherenkov Telescope Array Consortium 2019) will be able to see γ -ray sources an order of

* E-mail: mark.krumholz@anu.edu.au

magnitude fainter, and with an order of magnitude better resolution, than current instruments. The radio sky, and the galaxies that populate it, will be increasingly well revealed by the Square Kilometre Array (SKA) and its pathfinders. Finally, the upgrade of the South Polar IceCube neutrino telescope and the commissioning of the KM3NeT neutrino telescope will significantly improve our knowledge of the neutrino sky. This overall improvement in instrumentation will grant us the ability to directly probe the CR populations of star-forming galaxies with unprecedented depth and precision. This will, in turn, illuminate our understanding of CR transport.

However, there are a limited range of software tools capable of predicting the diverse signals that will be observable by the next generation of telescopes. Until very recently, CR-hydrodynamics simulations followed only a single CR energy bin; the most recent generation of simulations include a few distinct CR energies (e.g. Armillotta, Ostriker & Jiang 2021, 2022; Girichidis et al. 2022; Hopkins et al. 2022), but the computational cost of following multiple energies in the context of a fully self-consistent CR-hydrodynamics calculation means that these simulations achieve very limited spectral resolution in the observable signatures they predict, and rely on highly simplified treatments of the microphysical interaction between CRs and their environment (e.g. treating pp collisional losses as continuous, ignoring Klein–Nishina effects when calculating inverse Compton scattering). Fluid treatments that involve integration over broad bins in CR energy also necessarily have difficulty in treating sharp spectral features, for example sharp changes between streaming and diffusion at particular CR energies. Moreover, all of these methods have thus far proven too expensive to use in carrying out an extensive parameter study of how emission changes as one makes differing assumptions about the microphysics of CR interactions with the background plasma.

Conversely, a range of tools exist to predict CR observables using detailed microphysics, but generally only for highly simplified, time-independent background plasma states. The most prominent of these is GALPROP (Strong & Moskalenko 1998), but other examples include PICARD (Kissmann 2014), DRAGON (Evoli et al. 2017), and CRPROPA (Merten et al. 2017). These codes generally offer a much more detailed treatment of microphysics than is achieved in the CR-hydrodynamics simulations, and thus a correspondingly better prediction of observables, but at the price that they are not very flexible in terms of the assumed model of CR propagation, or in terms of the way that the background gas is described. Thus for example it would not be straightforward to use one of these tools to post-process a full 3D simulation and predict the observable, CR-driven emission from it. Nor for example can these codes easily handle a situation where the transport of CRs switches from primarily streaming to primarily diffusion as a function of CR energy.

This situation motivates us to introduce Cosmic Ray Interstellar Propagation Tool using Itô Calculus (CRIPTIC), a new software tool for the purpose of calculating CR transport and observable emission. CRIPTIC attempts to balance the advantages of the dedicated CR propagation codes – high-accuracy microphysics, high spectral resolution, accurate treatment of secondary particles, relatively high speed – with those of the CR-hydrodynamics codes – complex, multiphase gas backgrounds, with CR propagation properties that vary depending on the gas state. Of course it also has disadvantages, in that its complexity level and computational cost are higher than for GALPROP and similar software, and it does not achieve the full consistency between CR propagation and hydrodynamic evolution that comes from solving CRs and hydrodynamics together. None the less, CRIPTIC is unique in that it offers the ability to make realistic predictions for CR-driven emission at high spectral resolution

and high physical accuracy from 3D simulations of a complex, multiphase galactic gas ecosystem. In this regard the intended use of CRIPTIC for CRs is analogous to that of tools such as RADMC-3D (Dullemond 2012) or POWDERDAY (Narayanan et al. 2021) for photons: while these tools are too expensive to use in real time as part of a self-consistent radiation-hydrodynamics calculation, they offer much more realistic predictions of observable emission than would be possible using radiation-hydrodynamics simulations alone.

The remainder of this paper is laid out as follows. In Section 2, we describe the basic system of equations that CRIPTIC solves, and in Section 3 we describe the numerical method by which we solve them. We present validation tests in Section 4, and summarize and discuss future prospects in Section 5.

2 PHYSICAL MODEL

2.1 Formulation of the problem

CRIPTIC is intended to simulate the transport of CRs on scales significantly larger than the CR mean free path to pitch angle scattering. It therefore solves the Fokker–Planck Equation (FPE) for the evolution of the pitch angle-averaged CR distribution function $f(\mathbf{x}, p)$ as a function of position $\mathbf{x}$ and the (magnitude of the) CR momentum p ; future extensions may also include explicit evolution of the pitch angle distribution, but are beyond the scope of this paper. We solve a separate FPE for each species of CR tracked in a simulation. The equation we solve is (Skilling 1975; Zweibel 2017)

$$\begin{aligned} \frac{\partial f}{\partial t} = & \frac{\partial}{\partial x_i} \left(K_{ij} \frac{\partial f}{\partial x_j} \right) + \frac{1}{p^2} \frac{\partial}{\partial p} \left(p^2 K_{pp} \frac{\partial f}{\partial p} \right) \\ & - \frac{\partial}{\partial x_i} \left[\frac{\partial}{\partial p^3} (p^3 u_i) f \right] + \frac{\partial}{\partial p^3} \left[\frac{\partial}{\partial x_i} (p^3 u_i) f \right] \\ & - \frac{1}{p^2} \frac{\partial}{\partial p} (p^2 \dot{p}_{\text{cts}} f) - Lf + S, \end{aligned} \quad (1)$$

where we have added extra terms describing losses to the original transport equation derived by Skilling (1975); note that in this equation, although by assumption f depends only on the magnitude of CR momentum p , it still describes the number of CRs per unit volume in momentum space, i.e. the number of CRs with momentum from p to $p + dp$ is $4\pi p^2 f dp$ rather than $f dp$. Here and throughout Section 2, we adopt the convention that italic indices (i, j) go from 1 to 3 and Greek indices (α, β) go from 1 to 4, and we make use of the Einstein summation convention whereby repeated indices are understood to be summed over. The quantities appearing in this equation are the spatial diffusion tensor $\mathbf{K}$, the momentum diffusion coefficient¹ K_{pp} , the total (streaming plus advection) velocity of the CRs $\mathbf{u}$, the rate at which the magnitude of the CR momentum decreases due to continuous processes $\dot{p}_{\text{cts}}$ (those where the change in momentum per interaction is small), the catastrophic loss term L (describing processes that destroy CRs completely or cause large changes in their momenta), and the source function S . In general these quantities are functions of the properties of the background plasma at position $\mathbf{x}$, the CR momentum p , and the CR distribution f at that point in space – for example, the direction of CR streaming

¹Note that the subscript pp here indicates diffusion in momentum space, and should not be confused with proton–proton collisions; to avoid confusion, we will generically refer to proton interactions in general as nuclear inelastic collisions, so that pp is used only to indicate the momentum direction in phase space.

may depend on the gradient of f – in which case the problem is non-linear. The background plasma is threaded by a large-scale magnetic field $\mathbf{B}$, where large scale here means that $\mathbf{B}$ is averaged on scales much larger than the CR isotropization scale.

Before proceeding further, we pause to explain our formulation of the continuous and catastrophic loss terms. We approximate continuous loss processes as causing a continuous decrease in the momentum of each CR particle at a rate $\dot{p}_{\text{cts}}$. This corresponds to advection in momentum space at a velocity $-\dot{p}_{\text{cts}}\hat{r}_p$, where $\hat{r}_p$ is a radial (in momentum space) unit vector. The term we have included in the equation (1) is simply the divergence of $-\dot{p}_{\text{cts}}f\hat{r}_p$. By contrast, the term $-Lf$ we introduce for catastrophic losses corresponds to direct removal of CRs from the population, at a rate proportional to the number density of CRs at a given position; in equation (1), L is the probability per unit time of such a loss. If the loss process represents a discrete jump in momentum rather than complete destruction, there will also be corresponding source term $S \propto -Lf$. We describe in Section 2.3 the full set of microphysical processes we include and their mathematical representation in terms of continuous loss terms, catastrophic loss terms, and source terms.

2.2 From Fokker–Planck to stochastic differential equation

Our basic approach to solving equation (1) is to transform it from a PDE to a stochastic differential equation (SDE). This part of our treatment is similar to the methods proposed by Kopp et al. (2012) and Merten et al. (2017). To simplify the problem of transforming the FPE into an SDE, we will always work in a local coordinate system defined relative to the magnetic field. We therefore set up a TNB coordinate frame defined by the unit vectors

$$\hat{\mathbf{t}} = \frac{\mathbf{B}}{|\mathbf{B}|} \quad (2)$$

$$\hat{\mathbf{n}} = \frac{(\hat{\mathbf{t}} \cdot \nabla) \hat{\mathbf{t}}}{k} \quad (3)$$

$$\hat{\mathbf{b}} = \hat{\mathbf{t}} \times \hat{\mathbf{n}}. \quad (4)$$

Here, $\hat{\mathbf{t}}$, $\hat{\mathbf{n}}$, $\hat{\mathbf{b}}$, and k are the tangent vector, normal vector, binormal vector, and curvature of the local magnetic field. We therefore adopt the convention in equation (1) that indices $i = (1, 2, 3)$ correspond to the $(\hat{\mathbf{t}}, \hat{\mathbf{n}}, \hat{\mathbf{b}})$ directions, respectively. In our chosen frame the diffusion tensor $\mathbf{K}$ is diagonal, with elements

$$K_{ij} = \begin{cases} K_{\parallel}, & i = j = 1 \\ K_{\perp}, & i = j = 2 \text{ or } i = j = 3 \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

Note that we explicitly allow for the possibility of diffusion perpendicular to field lines, with diffusion coefficient $K_{\perp}$, and in this respect our equation differs from that given by Skilling (1975) or Zweibel (2017). Our reason for including this term is that the scale on which $\mathbf{B}$ is measured is *not* assumed to be smaller than the turbulent dissipation scale, and thus there may be turbulent fluctuations in the magnetic field on top of the large-scale guide field $\mathbf{B}$; this will often be the case in galactic or cosmological simulations, for example, where the simulation resolution is insufficient to resolve the magnetic dissipation scale. Thus we wish to leave open the possibility that, while CRs do not diffuse perpendicular to magnetic field lines (to leading order), the field lines themselves can wander perpendicular to the large-scale guide field, and this will induce an effective diffusion in the CRs perpendicular to the large-scale guide field (Beattie et al. 2022; Sampson et al. 2022).

In order to transform equation (1) into an SDE, we must first recast it in standard Fokker–Planck form

$$\frac{\partial f}{\partial t} = -\frac{\partial}{\partial q_{\alpha}} (A_{\alpha} f) + \frac{1}{2} \frac{\partial^2}{\partial q_{\alpha} \partial q_{\beta}} (D_{\alpha\beta} f), \quad (6)$$

where $\mathbf{A}$ is the drift vector, $\mathbf{D}$ is the diffusion tensor (which must be symmetric), and $\mathbf{q}$ is the vector of phase-space variables upon which f depends, which we take to be $\mathbf{q} = (x_1, x_2, x_3, p)$. In order to do so, we make a change of variables from $f(\mathbf{x}, p)$ to

$$\tilde{f}(\mathbf{x}, p) = 4\pi p^2 f(\mathbf{x}, p), \quad (7)$$

i.e. $\tilde{f}$ represents the probability density per radial distance in momentum space, while f is the probability density per unit volume in momentum space. The advantage of this change is that, unlike f itself, $\tilde{f}$ is invariant under a steady flow of CRs in momentum space.² With this change of variables, and making some further algebraic simplification, equation (1) becomes

$$\begin{aligned} \frac{\partial \tilde{f}}{\partial t} = & \frac{\partial^2}{\partial x_i \partial x_j} (K_{ij} \tilde{f}) + \frac{\partial^2}{\partial p^2} (K_{pp} \tilde{f}) \\ & - \frac{\partial}{\partial x_i} \left[\left(\frac{\partial K_{ij}}{\partial x_j} + u_i + \frac{p}{3} \frac{\partial u_i}{\partial p} \right) \tilde{f} \right] \\ & - \frac{\partial}{\partial p} \left[\left(\frac{\partial K_{pp}}{\partial p} + 2 \frac{K_{pp}}{p} - \frac{p}{3} \frac{\partial u_i}{\partial x_i} - \dot{p}_{\text{cts}} \right) \tilde{f} \right] \\ & - L \tilde{f} + \tilde{S}, \end{aligned} \quad (8)$$

where $\tilde{S} = 4\pi p^2 S$ is the rate per unit volume per unit linear momentum at which new CRs are injected.³ For computational purposes it is convenient to explicitly write out the total velocity $\mathbf{u} = \mathbf{v} + w\hat{\mathbf{t}}$, where $\mathbf{v}$ is the advection velocity of the background gas and $w\hat{\mathbf{t}}$ is the streaming velocity of the CRs along the field, in the frame comoving with the gas. Doing so, equation (8) becomes

$$\begin{aligned} \frac{\partial \tilde{f}}{\partial t} = & \frac{\partial^2}{\partial x_i \partial x_j} (K_{ij} \tilde{f}) + \frac{\partial^2}{\partial p^2} (K_{pp} \tilde{f}) \\ & - \frac{\partial}{\partial x_i} \left\{ \left[\frac{\partial K_{ij}}{\partial x_j} + v_i + \left(w + \frac{p}{3} \frac{\partial w}{\partial p} \right) \hat{t}_i \right] \tilde{f} \right\} \\ & - \frac{\partial}{\partial p} \left\{ \left[\frac{\partial K_{pp}}{\partial p} + 2 \frac{K_{pp}}{p} - \dot{p}_{\text{cts}} \right. \right. \\ & \quad \left. \left. - \frac{p}{3} \left(\frac{\partial v_i}{\partial x_i} + \frac{dw}{dx_i} \hat{t}_i + w \frac{d\hat{t}_i}{dx_i} \right) \right] \tilde{f} \right\} \\ & - L \tilde{f} + \tilde{S}. \end{aligned} \quad (9)$$

Ignoring the catastrophic loss and source terms for the moment, equation (9) corresponds to equation (6) with drift vector

$$A_{\alpha} = \begin{cases} \frac{\partial K_{\parallel}}{\partial x_{\alpha}} + v_{\alpha} + w + \frac{p}{3} \frac{\partial w}{\partial p}, & \alpha = 1 \\ \frac{\partial K_{\perp}}{\partial x_{\alpha}} + v_{\alpha}, & \alpha = 2 \text{ or } 3 \\ \frac{\partial K_{pp}}{\partial p} + 2 \frac{K_{pp}}{p} - \dot{p}_{\text{cts}} \\ \quad - \frac{p}{3} \left(\frac{\partial v_i}{\partial x_i} + \frac{\partial w}{\partial x_i} \hat{t}_i + w \frac{\partial \hat{t}_i}{\partial x_i} \right), & \alpha = 4 \end{cases} \quad (10)$$

²Note that this step is omitted in the derivation given by Kopp et al. (2012) and Merten et al. (2017), and as a result their equations for the momentum distribution are not correct. However, since none of the tests reported in their papers consider CR momentum evolution, the problem does not manifest in the published results.

³By contrast, S is the injection rate per unit volume in space per unit volume in momentum space.

Table 1. Summary of loss processes included. Entries give the process name and the types of particles affected, followed by the equation or expression we use for the various terms – $\dot{p}_{\text{cts}}$, L , ξ , ϕ , and $d\Psi/d\epsilon$ – that describe the rate of continuous momentum loss, the rate of catastrophic loss, the secondary multiplicity, the secondary momentum distribution, and the specific power radiated per CR primary, respectively. For secondary multiplicities and momentum distributions, we indicate the type of secondary particle in parentheses. A blank entry indicates that the process does not produce the indicated effect, e.g. nuclear inelastic scattering does not produce continuous momentum loss.

Loss process	Affects	$\dot{p}_{\text{cts}}$	L	ξ	$d\phi/dp$	$d\Psi/d\epsilon$
Nuclear inelastic scattering	p	–	Equation (17)	1 (p) $\sigma_{\pi^\pm}/\sigma_{\text{nuc}}(e^\pm)$	Equation (18) (p) Equation (21) (e^+ , e^-)	Equation (19)
Ionization	p, e^+ , e^-	Equation (22)	–	–	–	–
Coulomb	p, e^+ , e^-	Equation (33)	–	–	–	–
Synchrotron	e^+ , e^-	Equation (39)	–	–	–	Equation (40)
Bremsstrahlung	e^+ , e^-	Equation (48)	Equation (45)	1 (e^+ , e^-)	Equation (46)	Equation (49)
Inverse Compton	e^+ , e^-	Equation (59)	Equation (61)	1 (e^+ , e^-)	Equation (63)	Equation (64)
Positron annihilation	e^+	–	Equation (69)	–	–	Equation (72)

and diffusion tensor

$$D_{\alpha\beta} = \begin{cases} 2K_{\parallel}, & \alpha = \beta = 1 \\ 2K_{\perp}, & \alpha = \beta = 2 \text{ or } 3 \\ 2K_{\text{pp}}, & \alpha = \beta = 4 \\ 0, & \alpha \neq \beta \end{cases}. \quad (11)$$

Note that, in writing out the components of the drift vector and diffusion tensor, we have made use of the fact that K_{ij} is diagonal in our chosen coordinate system.

The Itô SDE corresponding to equation (8) is (section 6.1, Gardiner 2009)

$$dq_{\alpha}(t) = A_{\alpha}(\mathbf{q}, t) dt + d_{\alpha\beta}(\mathbf{q}, t) dW_{\beta}(t), \quad (12)$$

where

$$d_{\gamma\alpha}d_{\gamma\beta} = D_{\alpha\beta} \quad (13)$$

and $d\mathbf{W}(t)$ is a 4D Wiener process. Since $\mathbf{D}$ is diagonal in our chosen coordinate frame, we trivially have $d_{\alpha\beta} = \sqrt{D_{\alpha\beta}}$, so $\mathbf{d}$ is diagonal as well. For notational convenience we define the vector of diffusion coefficients $\kappa = (2K_{\parallel}, 2K_{\perp}, 2K_{\perp}, 2K_{\text{pp}})$, i.e. κ is just the vector of diagonal elements of $\mathbf{d}$, so the SDE becomes

$$dq_{\alpha}(t) = A_{\alpha}(\mathbf{q}, t) dt + \left(\sqrt{\kappa(\mathbf{q}, t)} d\mathbf{W} \right)_{\alpha}. \quad (14)$$

We refer to κ as the diffusion vector from this point forward, keeping in mind that it in fact just the vector of eigenvalues of a rank-2 tensor.

2.3 Microphysical processes: losses, secondaries, and observables

2.3.1 Formalism

The continuous and catastrophic loss terms $\dot{p}_{\text{cts}}$ and $L\tilde{f}$ can represent a range of microphysical interactions between CRs and their environment. In turn, the catastrophic loss terms generally have corresponding source terms $\tilde{S}$ that represent either new particles produced in the interaction, or existing CRs jumping discontinuously from a higher to a lower momentum; we will generically write these source terms in the form

$$\tilde{S}_s(p) = \sum_{s'} \int L_{s'}(p') \xi_{s,s'}(p') \left(\frac{d\phi}{dp} \right)_{s,s'}(p') \tilde{f}_{s'}(p') dp', \quad (15)$$

where $\tilde{S}_s(p)$ is the source function for members of species s with momentum p , the sum runs over all species s' , $L_{s'}(p')$ is the catastrophic loss rate for members of species s' with momentum p' , $\xi_{s,s'}(p')$ is the mean multiplicity for production of species s by the

loss of members of species s' with momentum p' ,⁴ $d\phi/dp_{s,s'}(p')$ is the distribution of momenta p for members of species s produced by the loss of a member of species s' with momentum p' (normalized so $\int (d\phi/dp)_{s,s'}(p') dp = 1$ for all p'), and $\tilde{f}_{s'}(p')$ is the CR distribution function for members of species s' evaluated at momentum p' . In the discussion that follows, we will characterize the source functions for secondaries in terms of their values for the multiplicity ξ and the momentum redistribution function $d\phi/dp$. We generically refer to members of species s' as primaries, even if they were themselves produced by another, earlier collision. Similarly, we refer to members of species s – whose appearance is described by $\tilde{S}_s$ – as secondaries, even if they are, in fact, just the initial CR particle after it has been scattered to lower energy.

In addition to computing the loss and source terms, it is also of interest to predict the observable emission from CRs. We write the rate of specific radiative emission per cosmic ray particle as $d\Psi/d\epsilon$, where ϵ is the photon energy. We also calculate the CR ionization rate via an analogous expression, since this is of interest for astrochemistry.

Different loss processes apply to different CR species, and neither our species list nor the set of processes we include are exhaustive; for example, at present CRIPIC does not treat heavy CR nuclei or the spallation losses they suffer. Both the particle and process list may be expanded in future releases. At present the code tracks CR protons, electrons, and positrons, and it includes accurate treatments of all significant loss processes for those species at energies from ≈ 1 MeV to ≈ 1 PeV propagating through typical ISM conditions; on the low-energy end this limit is imposed by adopting the relativistic limit when computing electron radiative losses (our treatment of collisional processes is valid down to ≈ 0.1 MeV), while on the high-energy side it is limited by the availability of tabulated or analytically approximated cross sections. Below we describe all the processes we include, and the methods we use to compute them; we summarize our final expressions for all processes in Table 1. Note that we do not require an explicit additional term to represent streaming losses, as is required in hydrodynamic treatments of CRs, because in the FPE such losses are automatically included in the $\nabla \cdot (w\hat{r})$ term in the drift vector.

The remainder of this section describes the various loss processes we include, and how we treat them. To help provide intuition for this discussion, we plot the various loss rates as a function of CR proton, electron, and positron energy for three example

⁴That is, $\xi_{s,s'}(p') = 2$ means that, on average, a loss of one particle of species s' with momentum p' leads to the production of two particles of species s .




Figure 1. Loss rates due to all of the processes included in CRIPTIC, as a function of CR kinetic energy T . For catastrophic losses (dashed lines) we show the catastrophic loss rate L , while for continuous processes (solid lines) the loss rate we show is computed as $L_{\text{cts}} = \dot{p}_{\text{cts}}/p$. Colours indicate different loss processes, as shown in the legend. The top row is for protons, the bottom for positrons; electrons are identical to positrons except that positron annihilation (brown dashed line) does not apply. Columns show, from left to right, loss rates in low-density ionized gas (H^+), moderate density atomic gas (H I), and dense molecular gas (H_2); see main text for exact properties of the gas in each of these regions.

environments in Fig. 1; here, the catastrophic loss rate is L , and we define the equivalent continuous loss rate as $\dot{p}_{\text{cts}}/p$ for a CR packet of momentum p . Our example environments are representative of (i) a diffuse ionized region (H^+ , first column), (ii) a medium density neutral atomic region (H I , second column), and (iii) a dense molecular region (H_2 , third column). These regions are characterized by, respectively, total densities of H nuclei $n_{\text{H}} = 10^{-3}$, 1, and 10^3 cm^{-3} , magnetic field strengths $B = 1$, 7, and $20 \mu\text{G}$, and dilute blackbody radiation fields, in addition to the cosmic microwave background, with temperature $T_{\text{BB}} = 5000$, 5000, and 10 K , and dilution factors $W_{\text{BB}} = 10^{-13}$, 10^{-13} , and 1; we explain T_{BB} and W_{BB} in more detail below. The H^+ region is composed of fully ionized H and singly ionized He with Asplund et al. (2009) protosolar abundances; if we define X_s as the number of members of a particular species s per H nucleon in the background gas, the H^+ region has $X_{\text{H}^+} = 1$, $X_{\text{He}^+} = 0.0955$, and $X_e = 1.0955$; the atomic region has $X_{\text{H}^0} = 0.99$, $X_{\text{H}^+} = 0.01$, $X_{\text{He}^0} = 0.0955$, and $X_e = 0.01$; and the molecular region has $X_{\text{H}_2} = 0.5$, $X_{\text{He}^0} = 0.0955$, and $X_e = 10^{-7}$.

2.3.2 Nuclear inelastic scattering

CR protons with kinetic energies above a threshold value $T_\pi = (2 + m_\pi/2m_p)m_\pi c^2$ can scatter inelastically off nuclei in the ISM, producing secondary particles, most commonly π mesons; here $m_p = 0.9383 \text{ GeV } c^{-2}$ and $m_\pi = 0.2797 \text{ GeV } c^{-2}$ are the rest masses of the proton and π^0 , respectively. In the process, the CR proton retains, on average, only a fraction $\eta_{\text{in}} \approx 1/2$ of its initial energy, which makes this a catastrophic loss process. We approximate the total cross-section for nuclear inelastic scattering using the analytical fitting formula proposed by Kafexhiu et al. (2014), which is well-calibrated against both collider experiments and particle physics simulations:

$$\sigma_{\text{nuc}} = \epsilon_{\text{nuc}} \sigma_{pp} \approx \epsilon_{\text{nuc}} \left[1 - \left(\frac{T}{T_\pi} \right)^{1.9} \right]^3 \times \left[30.7 - 0.96 \log \left(\frac{T}{T_\pi} \right) + 0.18 \log^2 \left(\frac{T}{T_\pi} \right) \right] \text{ mb}, \quad (16)$$

where σ_{pp} is the inelastic cross-section for pp collisions, $T = (\gamma - 1)m_p c^2$ is the kinetic energy of the incident proton in the frame of the ISM, $\gamma = (1 + p^2/m_p^2 c^4)^{1/2}$ is the proton Lorentz factor, and ϵ_{nuc} is the ‘nuclear enhancement factor’ that represents the increase in the cross-section relative to a pure H gas due to the presence of heavier nuclei. We compute the ϵ_{nuc} from Kafexhiu et al.’s equation (24), evaluated using the ISM elemental abundances provided by Draine (2011, their table 1.4), and for a beam of CR protons alone. Note that the cross-section drops to zero nearly discontinuously near $T = T_\pi$, as is clear from the top row of Fig. 1.

For protosolar abundances (Asplund et al. 2009), the mean gas mass per H nucleon is $\mu_{\text{H}} m_{\text{H}}$, where m_{H} is the hydrogen mass and $\mu_{\text{H}} = 1.4$, and thus we can write the catastrophic loss rate for nuclear inelastic collisions as

$$L_{\text{nuc}} = \sigma_{\text{nuc}} v \frac{\rho}{\mu_{\text{H}} m_{\text{H}}}, \quad (17)$$

where ρ is the total gas density, m_{H} is the hydrogen mass, and $v = c\sqrt{1 - \gamma^{-2}}$ is the proton velocity. Assuming each loss results in the CR proton surviving but having its kinetic energy reduced by a factor of $\eta_{\text{in}} = 2$, we can write the corresponding source function (equation 15) for ‘secondary’ protons as having a multiplicity $\xi_{p,p} = 1$ and a momentum redistribution function

$$\phi_{\text{nuc}, p \rightarrow p}(p|p') = \delta [T(p) - \eta_{\text{in}} T(p')], \quad (18)$$

where $T(p)$ is the kinetic energy of a proton with momentum p .

Inelastic collisions also produce both γ -rays (mainly via the decay of π^0 mesons) and secondary CR positrons and electrons (mainly via the decays of π^+ and π^- mesons). We compute the γ -ray emission using the parametrization provided by Kafexhiu et al. (2014), which gives the differential cross-section for production of photons of energy ϵ by CR protons of kinetic energy T . This yields an emitted specific power per CR proton

$$\frac{d\Psi}{d\epsilon} = \frac{f_{\text{nuc}}}{\epsilon} A_{\text{max}}(T) F(T, \epsilon) v \frac{\rho}{\mu_{\text{H}} m_{\text{H}}}, \quad (19)$$

where ϵ is the photon energy, and $A_{\max}(T)$ and $F(T, \epsilon)$ are parametrized functions of the CR kinetic energy T given by equations (11) and (12) of Kafexhiu et al. (2014).

To compute the source functions for CR electrons and positrons, we neglect the sub-dominant contribution from η and more massive mesons, and focus on that from $\pi^\pm$. We compute the cross-sections for production of these products following Yang, Kafexhiu & Aharonian (2018); near the production threshold, and $T < 2$ GeV, we use an interpolated cross-section taken from the tabulated values shown in Yang et al.'s fig. 4, while at larger energies we follow them in adopting⁵

$$\sigma_{\pi^\pm}(p) = \langle n_{\pi^\pm} \rangle \sigma_{\text{nuc}}, \quad \langle n_{\pi^\pm} \rangle = 0.78 w^{-1/4} (w - 2)^{3/4} - \frac{1}{6} \pm \frac{1}{3}. \quad (20)$$

Here, $\langle n_{\pi^\pm} \rangle$ is the energy-dependent multiplicity of production of $\pi^\pm$ and $w = \sqrt{s}/m_p c^2 = [2(1 + T/m_p c^2)]^{1/2}$ is the ratio of the centre-of-mass energy to the proton rest mass. Decays of $\pi^\pm$ then produce final state $e^\pm$, with the distribution of electron/positron momenta given by

$$\frac{d\phi_{\text{nuc}, p \rightarrow e^\pm}}{dp} \propto \frac{dE_{e^\pm}}{dp} \int f_e(x_{e^\pm}) f_\pi(x_{\pi^\pm}) dE_{\pi^\pm}. \quad (21)$$

Here, $x_{\pi^\pm} = E_{\pi^\pm}/E_p$ is the energy of the pion created in the collision normalized to the initial proton energy, $x_{e^\pm} = E_{e^\pm}/E_{\pi^\pm}$ is the energy of the electron or positron normalized to the pion energy, and $f_\pi(x)$ and $f_e(x)$ are the distributions of the normalized energies given by equations (6) and (36) of Kelner, Aharonian & Bugayov (2006), respectively. Thus our final expression for the $e^\pm$ source functions, in terms of equation (15), is that the multiplicity function $\xi_{e^\pm, p} = \sigma_{\pi^\pm}/\sigma_{\text{nuc}}$, and the momentum redistribution function is given by equation (21). Note that we do not at present follow final state neutrinos, although it would be a straightforward extension to the existing code to do so.

2.3.3 Ionization

At the energies with which we are concerned, ionization of the background gas is a continuous loss mechanism for all CR particles.⁶ We can generically write the resulting rate of momentum loss for a CR of species s' as

$$\dot{p}_{\text{cts}, \text{ion}} = \frac{v}{dT/dp} \frac{\rho}{\mu_H m_H} \sum_s X_s \mathcal{L}_{s', s, \text{ion}}, \quad (22)$$

⁵Note that in this expression we use σ_{nuc} , rather than fully adopting Yang et al.'s approach of using σ_{in} but then separately accounting for the contribution of heavier nuclei to pion production using the approach of Kafexhiu (2016); our approach amounts to assuming that heavier nuclei produce the same ratio of π^+ to π^- and the same pion energy distribution as H. The error associated with this approximation is negligible except at energies $\lesssim 0.1$ GeV, where other processes are generally more important in any event.

⁶Our assertion that electron ionization losses can be treated as continuous differs from the conclusions of Ivlev et al. (2021), who argue that ionization losses for electrons must be treated as catastrophic. The difference is the energy range of interest: while their computations follow electrons down to $\sim$ keV energies, we are limited to $\gtrsim 1$ MeV by our use of relativistic expressions for radiative loss rates. Calculations using the cross sections given below show that, for a 1 MeV electron, ionising collisions that change the energy of a CR electron by more than 10 per cent account for only 2.8×10^{-5} of all collisions, and are collectively responsible for only 8.9 per cent of the total energy loss rate; both these figures decrease as we go to higher electron energies. For this reason, we treat electron ionization losses as continuous.

where T and p are the kinetic energy and momentum of the CR particle undergoing loss, the sum runs over all species s with which the CR can interact, X_s is the abundance of that species per H nucleon, and $\mathcal{L}_{s', s}$ is the loss function for that species, given by

$$\mathcal{L}_{s', s, \text{ion}} = \int_0^{W_{s', s, \text{max}}} (W + I_s) \frac{d\sigma_{s', s, \text{ion}}}{dW} dW. \quad (23)$$

Here, I_s is the ionization potential of species s , W is the kinetic energy of the ejected electron, $W_{s', s, \text{max}}$ is the maximum ejected electron kinetic energy allowed by kinematics, and $d\sigma_{s', s, \text{ion}}/dW$ is the differential cross-section for ejection of electrons with kinetic energy W by collisions between CRs of species s' and targets of species s . We include H, H₂, and He as target species, since these overwhelmingly dominate the ionization losses; however, note that it is trivial to extend the formalism we describe below to include other targets.

We follow Ivlev et al. (2021) in taking our differential cross-sections for ionization by CR protons from the semi-analytical model of Rudd et al. (1992), which gives

$$\frac{d\sigma_{p, s, \text{ion}}}{dw} = \sigma_0 \left(\frac{R}{I_s} \right)^2 \frac{F_{1, s}(T_p) + F_{2, s}(T_p)w}{(1 + w)^3}, \quad (24)$$

where $w = W/I_s$, $\sigma_0 = 4\pi a_0^2 N_s$, a_0 is the Bohr radius, N_s is the number of electrons in the outer shell of the target species, $R = 13.6$ eV is the Rydberg energy, T_p is the proton kinetic energy, and $F_{1, s}$ and $F_{2, s}$ are empirical fitting functions given by equations (43) and (44) of Rudd et al., which depend on T_p and the properties of the target species, but not on W . The maximum energy allowed by kinematics is $W_{p, s, \text{max}} = 4(m_e/m_p)T_p - I_s$, and inserting this and equation (24) into equation (23), the loss function is

$$\mathcal{L}_{p, s, \text{ion}} = \sigma_0 I_s \left(\frac{R}{I_s} \right)^2 \left[F_{2, s}(t) \log(1 + w_{p, s, \text{max}}) + (F_{1, s}(t) - F_{2, s}(t)) \frac{w_{p, s, \text{max}}}{1 + w_{p, s, \text{max}}} \right], \quad (25)$$

where $w_{p, s, \text{max}} = 4t - 1$ and $t = (m_e/m_p)(T_p/I_s)$. The total ionization cross-section is

$$\sigma_{p, s, \text{ion}} = \sigma_0 \left(\frac{R}{I_s} \right)^2 \left[F_{2, s}(t) \frac{w_{p, s, \text{max}}}{2(1 + w_{p, s, \text{max}})^2} + F_{1, s}(t) \frac{w_{p, s, \text{max}}(2 + w_{p, s, \text{max}})}{2(1 + w_{p, s, \text{max}})^2} \right]. \quad (26)$$

Note from Fig. 1 that there is a clear inflection in the proton loss function near 1 GeV, where protons transition from sub-relativistic to relativistic.

Similarly, we take our differential cross sections for ejection of electrons of energy W by CR electrons of energy T_e from the RBEQ model of Kim, Santos & Parente (2000, their equation 19, modified as per the description of the BEQ approximation in their text),

$$\begin{aligned} \frac{d\sigma_{e, s, \text{ion}}}{dw} = & \sigma_0 \frac{\alpha^4}{2i'(\beta_t^2 + \beta_{u_s}^2 + \beta_{i_s}^2)} \times \\ & \left\{ \frac{Q_s - 2}{t + 1} \left(\frac{1}{1 + w} + \frac{1}{t - w} \right) \frac{1 + 2t'}{(1 + t'/2)^2} + \right. \\ & (2 - Q_s) \left(\frac{1}{(1 + w)^2} + \frac{1}{(t - w)^2} + \frac{i_s'^2}{(1 + t'/2)^2} \right) + \\ & \left. \left(\frac{Q_s}{(1 + w)^3} + \frac{Q_s}{(t - w)^3} \right) \left[\log \frac{\beta_t^2}{1 - \beta_t^2} - \beta_t^2 - \log(2i_s') \right] \right\}, \quad (27) \end{aligned}$$

where α is the fine structure constant and Q_s and U_s are the dimensionless dipole strength and mean outer shell electron kinetic energy for the target species; we take these latter two quantities from table 1 of Kim & Rudd (1994). In the expression above, primes indicate quantities normalized to the electron rest energy (i.e. $t' = T_e/m_e c^2$, $i'_s = I_s/m_e c^2$), lower case indicates quantities normalized to the ionization potential (i.e. $t = T_e/I_s$), and β 's indicate the β factor corresponding to a particular energy (i.e. $\beta_t = [1 - 1/(t' + 1)^2]^{1/2}$). The maximum ejected electron energy normalized to the ionization potential is $w_{e,s,\max} = (t - 1)/2$, and inserting this and equation (27) into equation (23) gives

$$\mathcal{L}_{e,s,\text{ion}} = I_s \sigma_0 f_{\text{ion}} (g_{\text{ion},1} + g_{\text{ion},2}) \quad (28)$$

where

$$f_{\text{ion}} = \frac{\alpha^4}{2i'_s (\beta_t^2 + \beta_{u_s}^2 + \beta_{i_s}^2)}, \quad (29)$$

$$\begin{aligned} g_{\text{ion},1} = & \frac{Q_s - 2}{t + 1} \left(\log \frac{2}{t + 1} + t \log \frac{2t}{t + 1} \right) \frac{1 + 2t'}{(1 + t'/2)^2} + \\ & (2 - Q_s) \left[\log \frac{(1 + t)^2}{4t} + \frac{i_s'^2}{(1 + t'/2)^2} \left(\frac{t - 1}{2} \right) \right] + \\ & Q_s \frac{(t - 1)^2}{2t(t + 1)} \left[\log \frac{\beta_t^2}{1 - \beta_t^2} - \beta_t^2 - \log(2i'_s) \right], \end{aligned} \quad (30)$$

$$\begin{aligned} g_{\text{ion},2} = & \frac{Q_s}{2} \left(1 - \frac{1}{t^2} \right) \left[\log \frac{\beta_t^2}{1 - \beta_t^2} - \beta_t^2 - \log(2i'_s) \right] + \\ & (2 - Q_s) \left[1 - \frac{1}{t} - \frac{\ln t}{t + 1} \frac{1 + 2t'}{(1 + t'/2)^2} + \frac{i_s'^2}{(1 + t'/2)^2} \frac{t - 1}{2} \right]. \end{aligned} \quad (31)$$

The total ionization cross-section is

$$\sigma_{e,s,\text{ion}} = \sigma_0 f_{\text{ion}} g_{\text{ion},2}. \quad (32)$$

At the $\gtrsim 1$ MeV energies with which we are concerned, positrons have almost exactly the same total ionization cross-sections as electrons (e.g. Knudsen et al. 1990), and it is therefore reasonable to assume similarly identical differential cross-sections. However, there is one subtlety: for electrons the kinematic limit $w_{e,s,\max} = (t - 1)/2$ is a result of the incident and ejected electrons being identical particles, so it is not possible to say which is the ‘primary’ CR electron and which is the ‘secondary’, ejected electron. The value of $w_{e,s,\max}$ for electrons amounts to treating whichever electron has lower energy as the secondary. For positrons, the incident and ejected particles are distinguishable, so it is not obvious what to choose for $w_{e^+,s,\max}$. Fortunately, this choice makes relatively little difference for $\gtrsim 1$ MeV energies, since losses are dominated by collisions that eject electrons with $w \ll w_{\max}$; we therefore adopt the same kinematic limit for positrons as for electrons, and thus the same loss function, but warn that this approach would not be valid at lower energies.

In addition to the calculating losses, CRIPTIC reports the total ionization rate for each target species. For a CR packet of species s' , the rate per primary CR at which ionizations of background species s occurs is $\zeta_s = \sigma_{s',s,\text{ion}} v_{\text{CR}}$, where v_{CR} is the CR velocity, and $\sigma_{s',s,\text{ion}}$ is the total ionization cross-section, given by equation (26) for protons and equation (32) for electrons and positrons.

2.3.4 Coulomb losses

In an ionized medium, all CRs lose energy via Coulomb interactions with the surrounding electrons. As with ionization, this process is well-approximated as continuous at the energies with which we are concerned. We take our loss rates from Gould (1972), interpolating smoothly between the expressions provided for the classical limit, $\beta < \alpha$ (where $\beta = v/c$ is the CR velocity normalized to c and α is the fine structure constant), the non-relativistic limit, $\alpha < \beta \ll 1$, and the ultrarelativistic limit, $1 - \beta \ll 1$; these transitions are visible in the loss rates plotted in Fig. 1, at least for protons. Our expressions are generically of the form

$$\dot{p}_{\text{cts,Coul}} = \left(\frac{dE_{\text{CR}}}{dp} \right)^{-1} \frac{e^2 \omega_p^2}{v} B_{\text{stop}}, \quad (33)$$

where E_{CR} is the CR energy,

$$\omega_p = \sqrt{\frac{4\pi n_e e^2}{m_e}} \quad (34)$$

is the plasma frequency, m_e is the electron mass, n_e is free electron density, and B_{stop} is the stopping number, which is a function of E_{CR} with a different functional form for different CR particle types. For protons, interpolating between Gould’s three cases gives

$$B_{p,\text{stop}} = \ln \left(\frac{2\gamma m_e c^2 \beta^2}{\hbar \omega_p} \right) + \frac{1}{2} \ln \left[1 + \left(\frac{\Gamma_C \beta}{\alpha} \right)^2 \right] - \frac{\beta^2}{2}, \quad (35)$$

where $\Gamma_C = 0.5615$ is a numerical constant. The equivalent expression for electrons is

$$\begin{aligned} B_{e^-,\text{stop}} = & \ln \left[\frac{\sqrt{2\delta(\gamma - 1)\beta m_e c^2}}{\hbar \omega_p} \right] + \frac{1}{2} \ln \left[1 + \left(\frac{\Gamma_C \beta}{\alpha} \right)^2 \right] \\ & + \frac{1}{2} \left[\left(1 + \frac{2\gamma - 1}{\gamma^2} \right) \ln(1 - \delta) + \frac{\delta}{1 - \delta} \right] \\ & + \frac{\delta^2}{4} \left(\frac{\gamma - 1}{\gamma} \right)^2, \end{aligned} \quad (36)$$

and for positrons is

$$\begin{aligned} B_{e^+,\text{stop}} = & \ln \left[\frac{\sqrt{2\delta(\gamma - 1)\beta m_e c^2}}{\hbar \omega_p} \right] + \frac{1}{2} \ln \left[1 + \left(\frac{\Gamma_C \beta}{\alpha} \right)^2 \right] \\ & + \frac{\delta^2}{4} \left(\frac{\gamma - 1}{\gamma + 1} \right)^2 \left[\frac{1}{2} + \frac{1}{\gamma} + \frac{3}{2\gamma^2} - \left(\frac{\gamma - 1}{\gamma} \right)^2 \left(\frac{2\delta}{3} - \frac{\delta^2}{2} \right) \right] \\ & + \delta \left[\frac{\delta}{8} \left(\frac{\gamma - 1}{\gamma} \right)^2 - \frac{\gamma^2 - 1}{2\gamma^2} \right] \\ & - \frac{\delta}{2} \left(\frac{\gamma - 1}{\gamma + 1} \right) \left[\frac{\gamma + 2}{\gamma} - \frac{\gamma^2 - 1}{\gamma^2} \delta + \left(\frac{\gamma - 1}{\gamma} \right)^2 \frac{\delta^2}{3} \right]. \end{aligned} \quad (37)$$

Here γ is the CR Lorentz factor, and δ is a fitting parameter for which we adopt the recommended value $\delta = 1/2$.

2.3.5 Synchrotron radiation

Synchrotron radiation is a continuous loss mechanism for CR electrons and positrons.⁷ The rate of energy loss is given by the

⁷In sufficiently strong magnetic fields synchrotron losses can be significant for protons as well, but at present we do not include these. It would be straightforward to do so in the future, however.

usual loss formula for an isotropic distribution of pitch angles (e.g. Blumenthal & Gould 1970),

$$\left(\frac{dE}{dt}\right)_{\text{sync}} = -\frac{4}{3}\sigma_T c\beta^2 \gamma^2 U_B, \quad (38)$$

where $\beta = v/c$, σ_T is the Thomson cross-section, and U_B is the magnetic energy density. We therefore have

$$\dot{p}_{\text{cts, sync}} = \frac{4}{3}\sigma_T c\beta^2 \gamma^2 U_B \left(\frac{dE_{e^\pm}}{dp}\right)^{-1}. \quad (39)$$

The corresponding time- and direction-averaged specific power that an observer sees per CR electron or positron is

$$\frac{d\Psi}{d\epsilon} = \frac{\sqrt{3}e^3 B h}{m_e c^2} \int_0^\pi \sin\alpha \frac{v}{v_{c,\perp}} \int_{v/(v_{c,\perp} \sin\alpha)}^\infty K_{5/3}(\xi) d\xi d\alpha, \quad (40)$$

where $v = \epsilon/h$ is the photon frequency and $v_{c,\perp} = 3eB\gamma^2/4\pi m_e c$ is the cut-off frequency for CRs with a pitch angle $\alpha = \pi/2$, and $K_{5/3}(\xi)$ is the modified Bessel function of order 5/3. Note that synchrotron radiation is not isotropic with respect to the local magnetic field; the quantity we compute here is the mean power radiated over 4π sr. Similarly, we do not at present compute the polarised intensity that would be seen by an observer in a particular direction.

2.3.6 Bremsstrahlung

Bremsstrahlung is a catastrophic loss process for electrons and positrons, since, at least in the relativistic regime, energy loss is dominated by photons whose energies are a significant fraction of the CR energy. We adopt the bremsstrahlung differential cross sections given by Blumenthal & Gould (1970),

$$\frac{d\sigma_{e^\pm, s, \text{br}}}{d\epsilon} = \frac{\alpha r_e^2}{\epsilon} \left\{ \left[1 + \left(\frac{E_{e^\pm} - \epsilon}{E_{e^\pm}} \right)^2 \right] \phi_1(\Delta) - \frac{2}{3} \frac{E_{e^\pm} - \epsilon}{E_{e^\pm}} \phi_2(\Delta) \right\} \quad (41)$$

where ϵ is the energy of the photon, $E_{e^\pm}$ is the total (rest plus kinetic) energy of the CR electron or positron, α is the fine structure constant, r_e is the classical electron radius,

$$\Delta = \frac{\epsilon m_e c^2}{4\alpha Z E_{e^\pm} (E_{e^\pm} - \epsilon)}, \quad (42)$$

is the screening factor, Z is the nuclear charge, and ϕ_1 and ϕ_2 are the screening functions. For species s consisting of unshielded charges (free protons, electrons, and He nuclei),

$$\phi_1 = \phi_2 = -4Z^2 \left[\ln(2\alpha Z \Delta) + \frac{1}{2} \right], \quad (43)$$

while for shielded nuclei (H, He, and He^+), we use the tabulated screening functions provided by Blumenthal and Gould.

The functional form of the differential cross-section requires some care in our numerical treatment; the most obvious approach is to integrate the differential cross-section to compute a total cross-section, and set the catastrophic loss rate to be proportional to it. The problem is that, as a result of the $1/\epsilon$ dependence, the total cross-section is logarithmically divergent, even though the total energy loss rate is finite. To handle this situation, we divide bremsstrahlung losses into a catastrophic component, representing losses due to photons with energies larger than $f_{\text{cat}} = 1/10$ of the CR kinetic energy, and a continuous part, accounting for losses due to lower energy photons.

For the catastrophic part the cross-section is

$$\sigma_{e^\pm, s, \text{br-cat}} = \int_{f_{\text{cat}} T_{e^\pm}}^{T_{e^\pm}} \frac{d\sigma_{e^\pm, s, \text{br}}}{d\epsilon} d\epsilon, \quad (44)$$

where $T_{e^\pm} = E_{e^\pm} - m_e c^2$ is the CR kinetic energy; the corresponding catastrophic loss rate

$$L_{\text{br}} = \frac{\rho}{\mu_H m_H} v \sum_s X_s \sigma_{e^\pm, s, \text{br-cat}}, \quad (45)$$

where v is the CR velocity, ρ is the gas density, and X_s is again the abundance of species s per H nucleon. The momentum distribution function for the scattered CRs is

$$\frac{d\phi_{\text{br}}}{dp} \propto \frac{dT_{e^\pm}}{dp} \sum_s X_s \left(\frac{d\sigma_{e^\pm, s, \text{br}}}{d\epsilon} \right)_{\epsilon=T_{e^\pm}(p')-T_{e^\pm}(p)}, \quad (46)$$

where the differential cross-section $d\sigma_{e^\pm, s, \text{br}}/d\epsilon$ is evaluated at a photon energy ϵ corresponding to the difference between the initial, $T_{e^\pm}(p')$, and final, $T_{e^\pm}(p)$, CR kinetic energies, and, following our catastrophic-continuous division, we set the catastrophic loss cross-section to zero if $T_{e^\pm}(p) < f T_{e^\pm}(p')$.

For the continuous part of the losses, we define a loss function by

$$\mathcal{L}_{e^\pm, s, \text{br}} = \int_0^{f_{\text{cat}} T_{e^\pm}} \epsilon \frac{d\sigma_{e^\pm, s, \text{br}}}{d\epsilon} d\epsilon, \quad (47)$$

analogously to the ionization loss function introduced in Section 2.3.3. The corresponding loss rate is

$$\dot{p}_{\text{cts, br}} = \frac{v}{dT_{e^\pm}/dp} \frac{\rho}{\mu_H m_H} \sum_s X_s \mathcal{L}_{s', s, \text{br}}. \quad (48)$$

We show the continuous and catastrophic loss rates in Fig. 1. Note that the continuous loss rate drops sharply at energies $T_{e^\pm} \gtrsim 0.1$ GeV for the H I- and H₂-dominated regions, but that a similar transition does not occur for the catastrophic loss rate. This is a direct result of the behaviour of atomic shielding, which is significant when $\Delta \ll 1$. Examining equation (42), we see that $\Delta \ll 1$ when $\alpha E_{e^\pm} \gg m_e c^2$ and $E_{e^\pm} \gg \epsilon$ are both satisfied. The former condition is met only when $T_{e^\pm} \gtrsim 0.1$ GeV, while the latter is met only for the continuous part of the loss rate, which is why we see shielding effects only for continuous losses at high energy. Also note that the bremsstrahlung loss rate, while it drops at $T_{e^\pm} \gtrsim 0.1$ GeV, eventually stabilizes and becomes constant at yet higher energy. This is as direct result of the ionization fraction being non-zero even in H I- and H₂-dominated regions; at sufficiently high energy, the continuous loss rate becomes dominated by the residual population of free protons and electrons, which are unaffected by shielding.

Finally, the specific power radiated by bremsstrahlung photons per CR primary is

$$\frac{d\Psi}{d\epsilon} = \frac{\rho}{\mu_H m_H} v \sum_s X_s \frac{d\sigma_{e^\pm, s, \text{br}}}{d\epsilon}. \quad (49)$$

Note that this includes photon emission at all energies; we separate catastrophic and continuous losses for the purposes of calculating CR propagation, but there is no need to separate them when computing photon emission.

2.3.7 Inverse Compton scattering

Inverse Compton (IC) scattering of electrons and positrons can be either a continuous or catastrophic loss process depending on the initial CR Lorentz factor γ and initial photon energy ϵ' . In the

Thomson limit, which applies when

$$\Gamma(\epsilon', \gamma) \equiv \frac{4\gamma\epsilon'}{m_e c^2} \ll 1, \quad (50)$$

continuous loss is a good approximation, while in the Klein–Nishina regime, $\Gamma \gtrsim 1$, the CR typically loses a substantial fraction of its energy with each scattering (Blumenthal & Gould 1970). CRIPTIC must be able to operate in both regimes, since, for example, we wish to be able to model both $\sim$ GeV CRs interacting with CMB photons ($\Gamma \sim 10^{-6}$) and $\sim$ 10 TeV CRs interacting with visible or UV photons ($\Gamma \sim 10^3$). We therefore make use of the general Klein–Nishina expression for the cross-section, rather than the simplified Thomson cross-section. In this general case, the differential rate at which IC scattering produces photons of energy ϵ is (Jones 1968; Blumenthal & Gould 1970)

$$\frac{d\dot{N}}{d\epsilon} = \frac{2\pi r_e^2 c}{\gamma^2} \int_0^\infty \frac{1}{\epsilon'} \frac{dN'}{d\epsilon'} S(\epsilon, \epsilon', \gamma) d\epsilon', \quad (51)$$

where r_e is the classical electron radius, $dN'/d\epsilon'$ is the specific number density of the photons being scattered, and

$$S(\epsilon, \epsilon', \gamma) = \begin{cases} 2q \ln q + (1+2q)(1-q) + \frac{1}{2} \frac{[\Gamma(\epsilon', \gamma)q]^2}{1+q} (1-q), & 0 < q < 1 \\ 0, & \text{otherwise} \end{cases} \quad (52)$$

where $q = \epsilon/[\Gamma(\epsilon', \gamma)(\gamma m_e c^2 - \epsilon)]$. Here S is a dimensionless function describing the shape of the Klein–Nishina cross-section.

To proceed further, we assume that the radiation field with which CRs are interacting can be described as a sum of dilute blackbodies, characterized by a temperature T_{BB} and a dilution factor W_{BB} ; the CMB has $W_{BB} = 1$, $T_{BB} = 2.73$ K, while the starlight field of the Milky Way is well-approximated by three components with $(W_{BB}, T_{BB}/K) = (7 \times 10^{-13}, 3000), (1.7 \times 10^{-13}, 4000), (1 \times 10^{-14}, 7500)$ (Mathis, Mezger & Panagia 1983; Draine 2011). The corresponding photon number density for each component is

$$\frac{dN'}{d\epsilon'} = W_{BB} \frac{8\pi}{(hc)^3} \frac{\epsilon'^2}{e^{\epsilon'/k_B T_{BB}} - 1}. \quad (53)$$

At this point it is convenient to define non-dimensional versions of the initial photon energy ϵ' , the final photon energy ϵ , and the blackbody temperature; we therefore define

$$x = \frac{\epsilon'}{k_B T_{BB}} \quad y = \frac{\epsilon}{\gamma m_e c^2} \quad \Gamma_{BB} = \frac{4\gamma k_B T_{BB}}{m_e c^2}. \quad (54)$$

With these definitions, we can rewrite equation (51) for a single component of the radiation field as

$$\frac{d\dot{N}}{dy} = \frac{\alpha^3}{8\pi\gamma^3} \left(\frac{c}{r_e}\right) W_{BB} \Gamma_{BB}^2 \frac{dF_{IC}}{dy}, \quad (55)$$

where α is the fine structure constant,

$$\frac{dF_{IC}}{dy} \equiv \int_{x_{\min}}^\infty \frac{x}{e^x - 1} S(y, x, \gamma) dx, \quad (56)$$

describes the differential scattering rate in normalized photon energy units, $x_{\min} = y/[\Gamma_{BB}(1-y)]$ is the minimum normalized initial photon energy that can produce a scattered photon with normalized energy y (i.e. the minimum value of x for which $q < 1$), and $S(y, x, \gamma)$ is given by equation (52), but with the substitution $\Gamma(\epsilon', \gamma) \rightarrow x\Gamma_{BB}$ and $q \rightarrow y/[\Gamma_{BB}(1-y)]$. The total scattering rate and energy loss rate due to a single radiation field component are then

$$\dot{N} = \frac{\alpha^3}{8\pi\gamma^3} \left(\frac{c}{r_e}\right) W_{BB} \Gamma_{e, BB}^2 \int_0^1 \frac{dF_{IC}}{dy} dy \quad (57)$$

$$\dot{E} = \frac{\alpha^3}{8\pi\gamma^2} \left(\frac{c}{r_e}\right) W_{BB} m_e c^2 \Gamma_{e, BB}^2 \int_0^1 y \frac{dF_{IC}}{dy} dy. \quad (58)$$

In the integrals above, the parts of the integrands at $y \ll 1$ are well-approximated as continuous loss, while those from y near unity correspond to losses that should be treated catastrophically. Following our approach with bremsstrahlung, we handle this by somewhat arbitrarily placing the boundary between the continuous at catastrophic regimes at $y = f_{\text{cat}} = 1/10$, and we therefore set the continuous loss rate to

$$\dot{p}_{\text{cts, IC}} = \left(\frac{dT_{e^\pm}}{dp}\right)^{-1} \frac{\alpha^3}{8\pi\gamma^2} \left(\frac{c}{r_e}\right) m_e c^2 \sum_i W_{i, BB} \Gamma_{i, BB}^2 G_{IC}(\Gamma_{i, BB}), \quad (59)$$

where we have defined

$$G_{IC}(\Gamma_{i, BB}) \equiv \int_0^{f_{\text{cat}}} y \frac{dF_{IC}}{dy} dy. \quad (60)$$

The sum in equation (59) runs over all the components of the radiation field, and $\Gamma_{i, BB}$ and $W_{i, BB}$ are the values that apply to the i th component. Similarly, the catastrophic loss rate is

$$L_{IC} = \frac{\alpha^3}{8\pi\gamma^3} \left(\frac{c}{r_e}\right) \sum_i W_{i, BB} \Gamma_{i, BB}^2 F_{IC}(\Gamma_{i, BB}), \quad (61)$$

where

$$F_{IC}(\Gamma_{i, BB}) \equiv \int_{f_{\text{cat}}}^1 \frac{dF_{IC}}{dy} dy. \quad (62)$$

In the limit $\Gamma_{BB} \rightarrow 0$, the functions F_{IC} and G_{IC} have the property that $G_{IC} \rightarrow (\pi^4/135)\Gamma_{BB}^2$ and $F_{IC}/G_{IC} \rightarrow 0$, so the continuous loss rate approaches the usual expression for the Thomson limit, and the catastrophic loss rate becomes negligible in comparison. On the other hand, in the limit $\Gamma_{BB} \rightarrow \infty$, we have $G_{IC}/F_{IC} \rightarrow 0$, and catastrophic losses dominate. This behaviour is visible in Fig. 1, where we see continuous losses being dominant at low CR energy and giving way to catastrophic losses at higher energy; also note that, in the H^+ and $H I$ regions, there are two distinct peaks of catastrophic loss, one at higher energy arising from the cosmic microwave background photon field with $T_{BB} = 2.73$ K and one at lower energy from the starlight field with $T_{BB} = 5000$ K.

For the part of IC losses that we treat as catastrophic, since electrons and positrons are conserved, the multiplicity $\xi_{e^\pm, e^\pm} = 1$. Consistent with our division between continuous and catastrophic losses, the momentum redistribution function is

$$\frac{d\phi_{IC}}{dp} = \frac{1}{E_{e^\pm}} \frac{dE_{e^\pm}}{dp} \sum_i \frac{W_{i, BB} \Gamma_{i, BB}^2 \Theta(y - f_{\text{cat}}) \frac{dF_{i, IC}}{dy}}{\sum_i W_{i, BB} \Gamma_{i, BB}^2 F_{IC}(\Gamma_{i, BB})}, \quad (63)$$

where $y = 1 - E_{e^\pm}(p')/E_{e^\pm}(p)$, and the purpose of the Heaviside step function $\Theta(y - f_{\text{cat}})$ is to enforce our approach that we only treat as catastrophic interactions that cause the CR energy to change by more than $f_{\text{cat}} = 10$ per cent at a time. Finally, the specific radiated power per CR primary is

$$\frac{d\Psi}{d\epsilon} = \frac{\alpha^3}{8\pi\gamma^3} \left(\frac{c}{r_e}\right) y \sum_i W_{i, BB} \Gamma_{i, BB}^2 \frac{dF_{i, IC}}{dy}, \quad (64)$$

where in this expression we do not include the step function because the observable emitted power includes both continuous and catastrophic losses.

2.3.8 Positron annihilation

The final (catastrophic) loss process we include is positron annihilation with electrons in the background gas. For the relativistic energies with which we are concerned, positronium formation is unimportant compared to direct annihilation, and Coulomb corrections to the annihilation cross-section are small (see the review by Prantzos et al. 2011), so the catastrophic loss rate is well-approximated by

$$L_{\text{annih}} = \sigma_d v \frac{\rho}{\mu_e m_H}, \quad (65)$$

where v is the CR positron velocity, $\mu_e m_H$ is the mean mass per electron ($\mu_e = 1.17$ for Asplund et al. 2009 abundances), and σ_d is the Dirac (1930) cross-section, which for a CR positron with Lorentz factor γ is given by

$$\sigma_d = \frac{\pi r_e^2}{\gamma + 1} \left[\frac{\gamma^2 + 4\gamma + 1}{\gamma^2 - 1} \ln(\gamma + \sqrt{\gamma^2 - 1}) - \frac{\gamma + 3}{\sqrt{\gamma^2 - 1}} \right]. \quad (66)$$

The energy dependence of the cross-section is visible in Fig. 1.

We obtain the specific radiated power per CR positron by first transforming to the centre-of-mass frame of the electron–positron collision, which has Lorentz factor $\sqrt{\gamma}$ relative to the lab frame defined by the background gas. In this frame, the electron and positron both have energy $E_{\text{CM}} = \sqrt{\gamma} m_e c^2$ and momentum $p_{\text{CM}} = (\gamma - 1)m_e c$, so each annihilation produces two photons with energy $\epsilon_{\text{CM}} = \sqrt{\gamma} m_e c^2$, and an angular distribution (equation 5.106, Peskin & Schroeder 1995)

$$\frac{dp}{d\mu_{\text{CM}}} \propto \frac{2 + \gamma + (\gamma - 1)\mu_{\text{CM}}^2 - 2[\gamma + (1 - \gamma)\mu_{\text{CM}}^2]^{-1}}{\gamma + (1 - \gamma)\mu_{\text{CM}}^2} \quad (67)$$

where μ_{CM} is the cosine of the angle between the direction of collision and the direction of the emitted photons. The corresponding angle measured in the lab frame is $\mu = (\mu_{\text{CM}} + \beta_{\text{CM}})/(1 + \beta_{\text{CM}}\mu_{\text{CM}})$, where $\beta_{\text{CM}} = \sqrt{1 - 1/\gamma}$ is the normalized velocity of the centre-of-mass frame, and the corresponding photon energy measured in the lab frame is

$$\epsilon = \frac{\epsilon_{\text{CM}}}{\sqrt{\gamma}(1 - \beta_{\text{CM}}\mu)} = \left[\gamma + \sqrt{\gamma(\gamma - 1)}\mu_{\text{CM}} \right] m_e c^2 \quad (68)$$

Thus the energy distribution of the photons produced must be proportional to $(d\mu_{\text{CM}}/d\epsilon)(dp/d\mu_{\text{CM}})$.

While this functional form describes the distribution of CR energies that would be measured by observers in the lab frame isotropically distributed around the direction of the collision, we must of course obtain the same energy distribution for the situation of interest to us, where a single observer at rest measures the photon energy distribution emitted by an isotropic collection of annihilating positrons averaged over all 4π sr. We can therefore write the quantity of interest to us, the rate per unit energy per CR positron emitted by an isotropic positron population as measured in the observer frame, as

$$\frac{d\dot{N}}{d\epsilon} = \frac{2L}{m_e c^2} \frac{dF_{\text{pos}}}{dy}, \quad (69)$$

where we have defined $y = \epsilon/m_e c^2$ as the ratio of photon energy to electron rest energy, the function dF_{pos}/dy describes the energy distribution of the received photons, and we adopt the normalization $\int (dF_{\text{pos}}/dy) dy = 1$, so that the total photon production rate integrated over all energies is $2L$, i.e. every positron that annihilates produces two photons. Since the energy distribution function $dF_{\text{pos}}/dy \propto (d\mu_{\text{CM}}/d\epsilon)(dp/d\mu_{\text{CM}})$, it is straightforward to

work out its functional form:

$$\frac{dF_{\text{pos}}}{dy} = \mathcal{N}^{-1} \left[\frac{y^4 - 4y^3 + 2\gamma(1 + 3\gamma)y^2 - 4\gamma^2(1 + \gamma)y + 2\gamma^2}{y^2(y - 2\gamma)^2} \right], \quad (70)$$

where dF_{pos}/dy is non-zero only for energies $y \in (y_-, y_+)$ with $y_{\pm} = \gamma \pm g$ and $g \equiv \sqrt{\gamma(\gamma - 1)}$. The normalization factor $\mathcal{N}$ required to ensure unit integral over this energy range is

$$\mathcal{N} = \frac{2}{\gamma} \left(3 - \frac{g}{\gamma} - 4\gamma + 4g \right) \cdot \left\{ 3g\gamma - g\gamma^2 - 4g\gamma^3 - \gamma [1 + \gamma(4\gamma^2 - \gamma - 4)] + [(3 - 10\gamma)\gamma + g - 6g\gamma + 6g\gamma^2 + 8g\gamma^3] \tanh^{-1} \frac{g}{\gamma} + 2\gamma^3(1 + 4\gamma) \sinh^{-1} \frac{g}{\sqrt{\gamma}} \right\}. \quad (71)$$

We can now write down our final expression for the specific power per CR primary,

$$\frac{d\Psi}{d\epsilon} = 2\sigma_d v \left(\frac{\rho}{\mu_e m_H} \right) y \frac{dF_{\text{pos}}}{dy}. \quad (72)$$

However, we caution that this expression only includes emission from positron losses ‘in flight’, which likely represent only a minority of total positron annihilations, with the balance occurring due to the formation of positronium after positrons have dropped to near-thermal energies via other loss processes (Prantzos et al. 2011). We do not include a treatment of positronium formation or the resulting emission in CRIPITIC, though it would be straightforward to apply such a model to the output of a CRIPITIC calculation, since CRIPITIC records the location and time at which each CR packet drops below the minimum momentum threshold at which we cease to follow it.

3 NUMERICAL METHOD

We can obtain solutions to the FPE, equation (8), by solving the corresponding SDE, equation (14), to obtain the trajectories through phase space $\mathbf{q}(t)$ for a large number of sample CR packets, including extra steps to account for losses and sources. The phase space distribution of those packets at any time t then provides an estimate of the phase space density $\tilde{f}(\mathbf{q})$ at that time. Each sample packet is characterized by a phase-space position $(\mathbf{x}, p)$, a weight Υ indicating the number of individual particles it represents, and the mass m and charge Ze of the particles that comprise it; each packet represents only a single species, but a computation may include an arbitrary number of species, each with its own distribution function $\tilde{f}$ and corresponding sample packets.

CRIPITIC advances the sample packets through a series of time-steps Δt . The procedure for updating from time $t^{(n)}$ at the end of the n th time-step to time $t^{(n+1)} = t^{(n)} + \Delta t^{(n)}$ has four parts, which we describe in detail in the subsequent sections:

- (i) If the diffusion vector, drift vector, or loss rate depends on the CR distribution function $\tilde{f}$ (i.e. if the problem is non-linear), estimate the required functions of $\tilde{f}$ at the position of each packet (Section 3.1).
- (ii) If any CR sources are present, inject new packets (Section 3.2).
- (iii) Advance all packets to time $t^{(n+1)}$ using an Euler-Maruyama (EM) update; in the process determine the next time step $\Delta t^{(n+1)}$ (Section 3.3).
- (iv) Check for production of secondaries during the time step, which are treated stochastically; if any secondary packets are

produced during the time step, update them to time $t^{(n+1)}$ as well, iterating until no packets remain (Section 3.4).

We first describe the operations of each of these steps in serial, and then how we modify the procedure for parallel computation in Section 3.5. We describe some general features of the CRIPTIC implementation in Section 3.6.

3.1 Step 1: reconstructing the distribution function

In a non-linear problem, the drift vector $\mathbf{A}$ or diffusion vector κ depend on the CR distribution function $\tilde{f}$ itself, so the first step in an advance is to reconstruct the distribution function seen by each sample packet so that these non-linear dependencies can be evaluated; in a linear problem we skip this step, as it is computationally expensive. As we discuss below, CRIPTIC is largely agnostic about the particular form of the non-linearity, and can accommodate a wide range of CR propagation models. However, it is not computationally practical to allow arbitrary functional dependence on $\tilde{f}$. We therefore limit the type of non-linearity we allow to what is by far the most common case. This, namely, is that the agent responsible for generating the non-linearity is resonant interactions between CRs and waves in the background plasma, and the waves are themselves generated by the CRs via the streaming instability. This constrains the functional form of the non-linear dependence on $\tilde{f}$, because CRs with a particular momentum p can only resonantly interact with waves whose wavelength is smaller than the CR gyroradius r_g . When dealing with multiple CR species, this condition is most conveniently expressed in terms of the CR rigidity $R = pc/|Ze|$, where Z is the CR charge in units of the elementary charge e . The gyroradius $r_g = R/B$, where B is the local magnetic field strength and, since this is fixed at any given position, the condition for resonant interaction then implies that CRs of rigidity R can experience non-linear interactions with other CRs whose rigidity satisfies $R' > R$. Given this consideration, we restrict CRIPTIC to computing non-linear effects that can be described in terms of a dependence of the propagation or loss rates for a CR of rigidity R only on integrals of $\tilde{f}$ over particles with rigidities $R' > R$. Other dependencies will be considered in future expansions.

With this physical picture in mind, CRIPTIC estimates the CR number density $n_{R'>R}$, pressure $P_{R'>R}$, and (kinetic) energy density $U_{R'>R}$, and the gradients of these quantities at the position of each packet, considering only the contributions from CRs with rigidity greater than or equal to that of the packet being considered; we will drop the $R' > R$ subscript from this point forward for brevity. For a packet with momentum p and charge Ze , these quantities are (Zweibel 2017)

$$\begin{pmatrix} n \\ P \\ U \end{pmatrix} = \sum_s \int_{pZ_s/Z}^{\infty} \begin{pmatrix} 1 \\ v_s p' \\ T_s \end{pmatrix} \tilde{f}_s dp' \quad (73)$$

where the sum runs over all CR species, $Z_s e$ is the charge on species s , $\tilde{f}_s$ is distribution function for species s evaluated at position $\mathbf{x}$, and v_s and T_s are the velocity and kinetic energy of a CR of species s with momentum p' ; note that these functions depend on species s because they depend on the particle mass. Thus n , P , and U are simply integrals of $\tilde{f}$ over momentum, evaluated with different weights – 1 for n , $v_s p$ for P , and T_s for U .⁸ The expressions for the gradients are completely analogous, simply replacing $\tilde{f}$ with $\nabla \tilde{f}$ in equation (73).

⁸Extensions of CRIPTIC to compute other quantities that are defined by similar weighted integrals over $\tilde{f}$ are trivial to implement if required, and simply require supplying the weight function for that quantity.

We evaluate these integrals by approximating them with Gaussian kernel density estimates; for each packet we define a bandwidth tensor $\mathbf{H}$ (computed as we describe below), and approximate the integrals above as

$$\begin{pmatrix} n \\ P \\ U \end{pmatrix} = \sum_{s,i} K_{\mathbf{H}}(\mathbf{x} - \mathbf{x}_{si}) \Upsilon_{si} \Theta_{si} \begin{pmatrix} 1 \\ v_{si} p_{si} \\ T_{si} \end{pmatrix}, \quad (74)$$

where the sum runs over all species s and all packets i belonging to that species, Υ_{si} , $\mathbf{x}_{si}$, p_{si} , v_{si} , and T_{si} are the weight, position, momentum, velocity, and kinetic energy of packet i of species s , Θ_{si} is unity for $p_{si} > p(Z_s/Z)$ and zero otherwise, and

$$K_{\mathbf{H}}(\mathbf{x}) = \sqrt{\frac{1}{8\pi^3 \det \mathbf{H}}} \exp\left(-\frac{1}{2} \mathbf{x}^T \mathbf{H}^{-1} \mathbf{x}\right) \quad (75)$$

is the usual 3D Gaussian kernel. The analogous expression for the gradients of these quantities are

$$\begin{pmatrix} \nabla n \\ \nabla P \\ \nabla U \end{pmatrix} = -\mathbf{H}_{\nabla}^{-1} \sum_{s,i} (\mathbf{x} - \mathbf{x}_{si}) K_{\mathbf{H}_{\nabla}}(\mathbf{x} - \mathbf{x}_{si}) \Upsilon_{si} \Theta_{si} \begin{pmatrix} 1 \\ v_{si} p_{si} \\ T_{si} \end{pmatrix}. \quad (76)$$

Note that the bandwidth $\mathbf{H}_{\nabla}$ used to estimate the gradient is not the same as that used to estimate the quantities themselves, as discussed below.

We evaluate the sums in equation (74) using an order $N \ln N$ algorithm based on a kd-tree decomposition. Our procedure is as follows. First, we sort the packets into a balanced kd-tree, and for each node in the tree we record the sum of the weights $\sum \Upsilon$ and squared weights $\sum \Upsilon^2$ for all packets contained in that node. Once the tree has been constructed, the next step is to determine the bandwidth tensor $\mathbf{H}$ for each packet. There is a vast body of literature on optimal methods for bandwidth selection, but the overriding constraint for us is that we require a method that operates quickly and without requiring global communication (for distributed memory calculations); the latter constraint rules out methods such as cross-validation or multistage plug-in selectors.

Instead, we make use of the tree structure itself to make an estimate of the local bandwidth, by choosing the bandwidth that brings a target number of neighbours $N_{\text{ngb, target}}$ within the kernel; our default value for this parameter is 1024, but users can choose alternate values. We define the effective neighbour number for each node of the tree as $N_{\text{ngb, node}} = (\sum \Upsilon)^2 / \sum \Upsilon^2$, and for each packet we start at the leaf of the tree that contains it, and climb the tree until we reach a node for which $N_{\text{ngb, node}} \geq N_{\text{ngb, target}}$ (or the root of the tree). At this point, we set the bandwidths $\mathbf{H}$ and $\mathbf{H}_{\nabla}$ for the packet by applying the optimal normal scale bandwidth selectors (equations 3.17 and 3.18, García-Portugués 2022),

$$\mathbf{H} = \left(\frac{4}{5}\right)^{2/7} N_{\text{ngb, node}}^{-2/7} \Sigma \quad (77)$$

$$\mathbf{H}_{\nabla} = \left(\frac{4}{7}\right)^{2/9} N_{\text{ngb, node}}^{-2/9} \Sigma, \quad (78)$$

where Σ is the covariance matrix for the points in the node.

Once we have selected a bandwidth for each packet, we use the tree to evaluate equations (74) and (76). We defer details of the algorithm to Appendix A.

3.2 Step 2: injecting packets

The second step in our algorithm is that each CR source present in the simulation volume adds new CR packets; in terms of equation (8), this represents the source term $\tilde{S}$. A source is characterized by its

position $\mathbf{x}_s$ at the start of the time-step, the first two derivatives of its position $\dot{\mathbf{x}}_s$ and $\ddot{\mathbf{x}}_s$, the species of CRs it produces, the total rate $\dot{N}$ at which it produces CRs (measured as particles injected per unit time), and the momentum distribution df/dp of those particles, where we normalize df/dp to have unit integral. Our current implementation uses sources with truncated power-law distributions in momentum $df/dp \propto p^q$ over the interval (p_0, p_1) , but extension to alternative functional forms of the momentum distribution is trivial.

When we inject packets, we assign each packet an injection time t_i drawn from a uniform distribution from $t^{(n)}$ to $t^{(n)} + \Delta t^{(n)}$, and an injection position $\mathbf{x}_i = \mathbf{x}_s + \dot{\mathbf{x}}_s(t_i - t^{(n)}) + \ddot{\mathbf{x}}_s(t_i - t^{(n)})^2/2$. Assigning packet momenta requires some subtlety, because the naive approach – drawing momenta from df/dp – is very computationally inefficient. For most realistic sources the momentum distribution is very steep, $df/dp \sim p^{-2.2}$, and so if every injected packet represents an equal fraction of the CR population, then a very large number of packets are needed to capture the behaviour at high momenta.

For this reason, we do not draw momenta from df/dp , but instead from an alternative distribution df_{samp}/dp that is generally flatter; our default choice is $df_{\text{samp}}/dp \propto p^{-1}$, corresponding to a uniform distribution in $\log p$, but users can alter this. To compensate for undersampling low- p packets compared to their true numbers, we increase the weight of those packets we do draw. Specifically, during a time step in which we draw a total of N_{packet} CR packets to represent the CRs injected by the given source, we set the weight of each packet we inject to

$$\Upsilon = \frac{\dot{N} \Delta t^{(n)}}{N_{\text{packet}}} \left(\frac{df/dp}{df_{\text{samp}}/dp} \right), \quad (79)$$

where df/dp and df_{samp}/dp are both evaluated at the momentum p of the newly drawn packet. The factor in parentheses ensures that the momentum distribution function of the injected packets, weighted by the packet weights Υ , follows df/dp .

The number of packets N_{packet} injected by each source is set by a user-specified packet injection rate $\dot{N}_{\text{packet}}$, which specifies the number of primary CR packets injected per unit time by all sources in the calculation; this choice determines the trade-off between computational cost and fidelity in sampling the distribution function, and the optimal choice is necessarily problem-dependent. When only a single source is present, we trivially have $N_{\text{packet}} = \dot{N} dt$. In the more general case where multiple sources are present, we assign each source a weight

$$\Upsilon_s = \dot{N} \int \frac{df_{\text{samp}}}{dp} dp, \quad (80)$$

and then

$$N_{\text{packet},i} = \frac{\Upsilon_{s,i}}{\sum_j \Upsilon_{s,j}} \dot{N}_{\text{packet}} \Delta t^{(n)} \quad (81)$$

packets for source i , where the sum runs over all sources present. This ensures that the total number of packets injected is $\dot{N}_{\text{packet}} dt$, and that the total momentum distribution of all packets injected is distributed as df_{samp}/dp .

Once packets have been injected, we reconstruct the number density n , pressure P , and energy density U at their phase space locations using the procedure described in Section 3.1, exactly as for the packets that already exist at the start of the time-step.

3.3 Step 3: advancing packets

Consider a sample CR packet that starts a time step with phase space position $\mathbf{q}^{(*)}$ and weight $\Upsilon^{(*)}$, and let t^* be the time at which the

packet starts the step; for packets that existed at the start of the time step $t^* = t^{(n)}$ and similarly for $\mathbf{q}^{(*)}$ and weight $\Upsilon^{(*)}$, while for newly created packets $t^* = t_i$, where t_i is the time at which that packet was injected, and the phase space position $\mathbf{q}$ and weight Υ correspond to those with which the packet was injected. We must advance the packet to time $t^{(n+1)}$ through a series of sub-steps. We begin each sub-step by computing the drift and diffusion vectors $\mathbf{A}$ and κ , and the sum of the catastrophic loss rates due to all processes L , given the current properties and phase space position of the packet; if the distribution function $\tilde{f}$ or quantities derived from it are required, we use the reconstructed value obtained in Step 1. From these, we compute a series of time step constraints associated with spatial drift, spatial diffusion, momentum drift, momentum diffusion, and catastrophic loss:

$$\Delta t_{\text{x-drift}} = \frac{\Delta x}{\sqrt{A_1^2 + A_2^2 + A_3^2}} \quad (82)$$

$$\Delta t_{\text{x-diff}} = \frac{\Delta x^2}{\max(K_{\parallel}, K_{\perp})} \quad (83)$$

$$\Delta t_{\text{p-drift}} = \frac{p}{|A_4|} \quad (84)$$

$$\Delta t_{\text{p-diff}} = \frac{p^2}{K_{\text{pp}}} \quad (85)$$

$$\Delta t_{\text{loss}} = \frac{1}{L}. \quad (86)$$

Here Δx is the smaller of (i) a user-specified length-scale that describes the structure of the background gas through which the CR packets move (and thus the size scale on which $\mathbf{A}$ and κ might be expected to vary) and (ii) the square root of the geometric mean of the eigenvalues of the bandwidth tensor $\mathbf{H}$, which describes the typical size of the kernel used to estimate $\tilde{f}$. We then set the overall time step to

$$\Delta t = \min \left(t^{(n+1)} - t^{(*)}, \frac{C}{\Delta t_{\text{x-drift}}^{-1} + \Delta t_{\text{x-diff}}^{-1} + \Delta t_{\text{p-drift}}^{-1} + \Delta t_{\text{p-diff}}^{-1} + \Delta t_{\text{loss}}^{-1}} \right), \quad (87)$$

where C is a user-specified tolerance with a default value of 0.25; note that the sub-step size $\Delta t \leq \Delta t^{(n)}$.

Once the time step has been set, we first update the packet weight via

$$\Upsilon^{(\dagger)} = \Upsilon^{(*)} e^{-L \Delta t}, \quad (88)$$

where $(\dagger)$ indicates the state after the update, and we next update the momentum through a standard EM step (Gardiner 2009),

$$p^{(\dagger)} = p^{(*)} + A_4 \Delta t + \eta \sqrt{K_{\text{pp}} \Delta t}, \quad (89)$$

where η is a random variable drawn from a distribution with unit mean and zero variance. At this point we check for creation of secondaries, a procedure we describe in Section 3.4, and we delete packets for which p or Υ/Υ_i , where Υ_i is the packet weight at injection, fall below user-specified tolerances; this is to limit the use of computational resources following packets that represent either a negligible number of particles or that have fallen to energies below those in which we are interested.

The final step is to compute the new spatial position, which we also do through an EM update,

$$x_i^{(\dagger)} = x_i^{(*)} + A_i \Delta t + \eta_i \sqrt{\kappa_i \Delta t}, \quad (90)$$

where η is a vector of three independent random variables with unit mean and zero variance. However, there is a subtle difficulty

in implementing this expression: the positions here are written in the local TNB basis, which is not the same at every position in space. In principle we could handle this issue simply by evaluating the TNB basis vectors in our fixed coordinate system at the start of every sub-step, using the equation above to compute the displacement of the CR packet along these basis vectors, and then transforming these displacements into our fixed coordinate system. However, due to the first-order nature of the EM scheme,⁹ this approach leads to considerable numerical drift of packets across field lines. We therefore follow Merten et al. (2017) in circumventing this problem by dividing each EM step into a series of sub-steps during each of which we recompute the TNB basis vectors, carrying out those sub-steps using an adaptive, higher order method that provides vastly better field-line tracing.

Specifically, let ξ be the position in the fixed frame of the simulations, and define an effective ‘velocity’ in the local TNB frame by

$$\dot{x}_i = A_i^{(v=0)} \Delta t + \eta_i \sqrt{\frac{k_i}{\Delta t}}, \quad (91)$$

where the superscript ($v = 0$) indicates the drift vector $\mathbf{A}$ evaluated with the advection velocity $\mathbf{v}$ set to zero; we separate out the advection velocity because it is *not* defined relative to the TNB basis. We then carry out a Runge–Kutta–Fehlberg 4th–5th order (RKF45; Fehlberg 1970) update of the position in the fixed frame ξ by setting the derivative at each stage of the RKF45 update to

$$\dot{\xi} = \dot{x}_1 \hat{\mathbf{t}} + \dot{x}_2 \hat{\mathbf{n}} + \dot{x}_3 \hat{\mathbf{b}} + \mathbf{v}, \quad (92)$$

where the basis vectors $\hat{\mathbf{t}}$, $\hat{\mathbf{n}}$, and $\hat{\mathbf{b}}$ are recomputed at every stage, taking into account how the TNB basis vectors change as the packet moves. If the field does not vary in space, then the basis does not change, and equation (92) reduces to the standard EM update. As is usual with the RKF45 update, at the same time we compute the update, we can also compute an error estimate; if that error estimate exceeds some specified tolerance, we divide the RKF45 time step Δt into two sub-steps of size $\Delta t/2$, repeating the subdivision recursively as necessary until the error estimate drops below a specified tolerance. In this way we ensure that packets follow field lines accurately.

At this point, we have updated all packet variables for the sub-step Δt . We repeat this process until each packet is advanced to the final time $t^{(n+1)}$.

3.4 Step 4: catastrophic losses and secondaries

One of the steps in the advance procedure described in Section 3.3 is to create secondary packets, where secondary here is used to mean any packet created from another packet, rather than from a source or present in the initial conditions. As described in Section 2.3, the source function describing the secondaries created by some catastrophic loss process i can be characterized in terms of the loss rate L_i for the process, the multiplicity function $\xi_{i,s}$ that describes the multiplicity of secondaries of species s for that process, and the momentum distribution function $\phi_{i,s}$ for those secondaries. Note that L_i , $\xi_{i,s}$, and $\phi_{i,s}$ are all functions of the momentum p' of the CRs undergoing loss, but in the discussion below we do not write out this functional dependence explicitly for compactness.

⁹Higher order schemes such as the Milstein algorithm (Gardiner 2009) are unfortunately not usable for our problem, due to the potentially complex dependence of the diffusion coefficients on position.

Now consider a time step of length Δt during which a CR packet of statistical weight Υ experiences a total loss rate from all processes $L = \sum_i L_i$. As a result, the weight of the primary packet is reduced to $\Upsilon e^{-L\Delta t}$. The total number of secondaries of species s created by loss process i in the course of this evolution is

$$\Upsilon_{i,s,\text{sec}} = \xi_{i,s} \frac{L_i}{L} (1 - e^{-L\Delta t}) \Upsilon. \quad (93)$$

A naive implementation of secondary production would be to create secondaries with this weight every time step. However, doing so would quickly make the calculation impossibly expensive due to the rapid proliferation of packets. We therefore instead set a probability $p_{i,s,\text{sec}} \leq 1$ that each packet spawns a secondary of species s via loss process i during each sub-step of its advance (Section 3.3), and increase the statistical weight of the secondaries by a factor $1/p_{i,s,\text{sec}}$ to compensate, so that the expected statistical weight has the correct value. To be precise, for any secondaries we do create, we assign them a weight

$$\Upsilon_{i,s} = \frac{\xi_{i,s}}{p_{i,s,\text{sec}}} \frac{L_i}{L} (1 - e^{-L\Delta t}) \Upsilon \quad (94)$$

If a secondary is created, we assign its initial position to be the same as that of its parent at the start of the sub-step, and we assign its initial momentum by drawing from the momentum redistribution function $\phi_{i,s}$.

We set the secondary creation probability $p_{i,s,\text{sec}}$ for each process and species to a value proportional to the expected secondary creation rate $\xi_{i,s} L_i$. We then normalise the probability as follows:

$$p_{i,s,\text{sec}} = f_{\text{sec}} \frac{L_i \xi_{i,s}}{\sum_j L_j \sum_{s'} \xi_{j,s'}} (1 - e^{-L\Delta t}), \quad (95)$$

where the sums in the denominator run over all processes j and secondary species s' , and f_{sec} is a user-settable dimensionless parameter that functions analogously to $\dot{N}_{\text{packet}}$ in that it parametrizes the trade-off between fidelity and computational cost in tracking secondaries. With this choice, the weight assigned to each secondary created then reduces to

$$\Upsilon_{i,s} = \frac{\sum_j L_j \sum_{s'} \xi_{j,s'}}{f_{\text{sec}} L} \Upsilon. \quad (96)$$

To understand the meaning of the parameter f_{sec} , first note that the expected number of secondaries created during a time-step, summing over all loss processes and all species, is

$$\langle N_{\text{sec}} \rangle_1 = \sum_{i,s} p_{i,s,\text{sec}} = f_{\text{sec}} (1 - e^{-L\Delta t}). \quad (97)$$

After N such steps the expected number of secondaries created is $\langle N_{\text{sec}} \rangle = N \langle N_{\text{sec}} \rangle_1$, and the weight of the original packet will have been reduced by a factor $e^{-NL\Delta t}$. Thus if we let Υ_i and Υ_f be the initial and final weights of the packet undergoing losses, and adopt the limit of small time steps, $L\Delta t \ll 1$, then we can write the expected number of secondaries created as

$$\langle N_{\text{sec}} \rangle = f_{\text{sec}} \ln \frac{\Upsilon_i}{\Upsilon_f}. \quad (98)$$

Thus we see that f_{sec} determines the mean number of secondary packets created per e -folding of the primary packet weight; a value of $f_{\text{sec}} = 1$ corresponds to creating an average of one secondary per e -folding, while a value of $f_{\text{sec}} = 1/\ln 10 \approx 0.434$ corresponds to creating one secondary per factor of 10 reduction in the weight of the primary. We choose the latter as our default, but the optimal choice likely varies from problem to problem, depending on how much loss the primary particles suffer.

3.5 Parallelism

CRIPITIC is parallelized using a hybrid MPI-openMP model. Within a single MPI rank, it uses openMP threads to carry out the advance for all the packets owned by that rank. This involves two synchronization points: one at the end of the step when we compute the next time step, and one after calculating any non-linear terms in the drift and diffusion vectors prior, since this calculation must be completed prior to moving any packets.

The MPI parallelism has two parts. The first is a step to distribute the packets across ranks to maintain load balance. For this purpose we use the PANDA algorithm described by Patwary et al. (2016); briefly summarizing here, the algorithm constructs a global kd-tree and uses it to partition packets between ranks based on their spatial position. Construction of a kd-tree requires two steps at each level of partition: choosing a dimension along which to divide the packets, and then choosing a point in that dimension to partition the data. The first step is straightforward to carry out in parallel: the ranks compute the variance of positions in each dimension, and since this involves only summation operations, parallelizing is straightforward. We choose to split along the dimension of maximum variance. The second step, choosing a partition point, requires more care, since it is expensive to find the median point – the traditional choice for the partition – in a distributed memory parallel calculation. Instead, we find the median approximately by selecting a set of sample points from all ranks to mark the edges of a histogram, and counting how many points on each rank fall into the various histogram bins. The counts can then be reduced in parallel, allowing us to approximate the value of the median. This process repeats recursively at each level of the tree, until the desired level of leaves are created. Each leaf is assigned to an MPI rank, and the packets within each leaf are then sent to that rank. In practice, since the tree changes little from time step to time step, during most time steps the boundaries of leaves move little, and thus little data need be communicated to maintain load balance. Once the leaves of the global tree have been assigned to MPI ranks, we construct local kd-trees below those leaves, exactly as we do in a non-parallel calculation.

The second part of MPI parallelism is to carry out kernel density estimation in parallel when packets are distributed across ranks. As with the remainder of the tree algorithm we use for the kernel density computation, we defer details to Appendix A.

3.6 Implementation notes

CRIPITIC is written in C++, based on the C++17 standard; we have avoided C++ 20 features to ensure compatibility with somewhat older compilers. A key aspect of the design is to maximize user flexibility in specifying (1) the initial conditions, (2) the properties of the background gas and radiation field through which CRs propagate, and (3) the underlying plasma physical model that describes CR propagation. While flexibility in initial conditions is standard in simulation codes, flexibility in the background gas state and propagation model is more challenging.

In CRIPITIC, we achieve this using C++ classes. The state of the background gas is defined by a pure virtual interface class, which can be specialized by a user to describe an arbitrary gas distribution. We provide specializations for some standard cases, for example where the gas distribution is specified in terms of an analytical function, a static Cartesian grid, or a series of snapshots in time that are each stored on Cartesian grids but, given this flexibility, users can implement their own classes to describe arbitrary time- and

position-dependent magnetic fields, gas densities, ionization states, compositions, and background radiation fields.

We take a similar approach to CR propagation. In practice, CR propagation in CRIPITIC is defined using a pure virtual interface class `Propagation`, which defines the call operator `Propagation::()` as a pure virtual function that takes as input the spatial position $\mathbf{x}$, the time t , the properties of the CR packet (type of particle, momentum, etc.), the properties of the background gas (total density, ion density, composition, magnetic field, etc.), and the CR field quantities n , P , and U and their gradients. This function must return all of the quantities that appear in the drift vector $\mathbf{A}$ (equation 10) and the diffusion tensor $\mathbf{D}$ (equation 11): the parallel diffusion coefficient $K_{\parallel}$ and its spatial gradient $\nabla K_{\parallel}$, the perpendicular diffusion coefficient $K_{\perp}$ and its spatial gradient $\nabla K_{\perp}$, the momentum diffusion coefficient K_{pp} and its derivative with respect to momentum $\partial K_{pp}/\partial p$, and the streaming speed w and its spatial gradient ∇w and derivative with respect to momentum $\partial w/\partial p$.

To define a CR propagation model, the user defines a class derived from `Propagation` that provides an implementation of the call operator and computes the required outputs from the provided inputs; the implementation of this function is entirely up to the user, and thus, for example, CR propagation can include arbitrary combinations of streaming and diffusion, which can depend in arbitrary ways on position, time, CR properties, gas properties, and the field quantities n , P , and U and their gradients. As with the gas properties, we provide implementations for some standard cases – for example, a model where the CR diffusion coefficient is a power-law function of CR momentum, and where CRs stream down field lines at the ion Alfvén speed – but users are not limited to these choices. The only restrictions are that, in its current form, CRIPITIC cannot capture CR propagation coefficients that depend on something other than the provided inputs listed above, or where CR propagation is not describable by the pitch angle-averaged Fokker–Planck equation (e.g. because the pitch angle distribution is not close to isotropic).

One implication of this flexibility is that CRIPITIC can be run using exactly the same interstellar gas and radiation field distributions, and CR propagation models, as standard CR propagation codes such as GALPROP (Strong, Moskalenko & Ptuskin 2007).

4 CODE TESTS

Here, we describe the various validation tests to which we have subjected CRIPITIC.

4.1 Transport tests

Our first batch of tests evaluates CRIPITIC’s performance in simulating CR transport, including the step of reconstructing the CR field where necessary because transport rates depend non-linearly on it. For all the tests in this section, we disable all catastrophic and continuous loss terms, so we are testing the transport parts of the code only. In these tests we will characterize the performance of the code in terms of its L^1 error; we do not experiment with varying the number of packets explicitly, but below we show that the errors we obtain are generally consistent with Poisson noise, and thus in general we expect the error to depend on number of packets used in a given simulation as $L^1_{\text{err}} \propto N_{\text{packet}}^{-1/2}$.

4.1.1 Anisotropic diffusion

Our first test is to validate CRIPITIC’s treatment of diffusion, including anisotropy and momentum-dependence of the diffusion coefficients.

We consider a uniform region containing gas at rest, threaded by a uniform magnetic field in the $\hat{x}$ direction. The CR diffusion coefficients parallel and perpendicular to the field have a fixed ratio χ and vary as power laws in the CR momentum, i.e.

$$K_{(\parallel,\perp)} = (\chi, 1) \cdot K_0 \left(\frac{p}{p_0} \right)^q, \quad (99)$$

where χ, K_0, p_0 , and q are constants. There is no momentum diffusion or streaming. The domain is initially empty of CRs, but at $t = 0$ a point source of CR protons turns on at the origin. The source is monochromatic, and is characterized by its luminosity $\mathcal{L}$ and by the momentum p_0 of the protons it injects. For this problem setup, the Fokker–Planck equation reduces to

$$\frac{\partial \tilde{f}}{\partial t} = K \left(\chi \frac{\partial^2 \tilde{f}}{\partial x^2} + \frac{\partial^2 \tilde{f}}{\partial y^2} + \frac{\partial^2 \tilde{f}}{\partial z^2} \right) + \frac{\mathcal{L}}{T} \delta(\mathbf{r}) \delta(p - p_0) \Theta(t), \quad (100)$$

where $\mathbf{r} = (x, y, z)$, $K \equiv K_0(p/m_p c)^q$, $\Theta(x)$ is the Heaviside step function, and

$$T = m_p c^2 \left[\sqrt{1 + \left(\frac{p_0}{m_p c} \right)^2} - 1 \right] \quad (101)$$

is the kinetic energy of the injected protons.

Making the change of variable $x = \sqrt{\chi} x'$ reduces this problem to a constant-coefficient diffusion equation

$$\frac{\partial \tilde{f}}{\partial t} = K \left(\frac{\partial^2 \tilde{f}}{\partial x'^2} + \frac{\partial^2 \tilde{f}}{\partial y^2} + \frac{\partial^2 \tilde{f}}{\partial z^2} \right) + \frac{1}{\sqrt{\chi}} \frac{\mathcal{L}}{T} \delta(\mathbf{r}') \delta(p - p_0) \Theta(t), \quad (102)$$

where $\mathbf{r}' = (x', y, z)$; note the extra factor of $1/\sqrt{\chi}$ in the final term, which arises from the change of variable. The Green's function for the spatial distribution of the CRs is then

$$G(r', t) = \frac{1}{(4\pi K t)^{3/2}} e^{-r'^2/4Kt}, \quad (103)$$

where $r' = |\mathbf{r}'|$, so for the case our case of a source that turns on at $t = 0$, the solution for the spatial distribution for $t > 0$ is

$$\begin{aligned} \tilde{f}(r', t) &= \int_0^t \frac{1}{\sqrt{\chi}} \frac{\mathcal{L}}{T} \delta(p - p_0) G(r', t') dt' \\ &= \frac{1}{\sqrt{\chi}} \frac{\mathcal{L}}{T} \frac{1}{4\pi K r'} \operatorname{erfc} \left(\frac{r'}{2\sqrt{Kt}} \right) \delta(p - p_0). \end{aligned} \quad (104)$$

The corresponding CR energy density is simply $U_{\text{CR}} = T \int f(r', t) dp$.

We test CRIPTIC by simulating this problem using $K_{\perp} = 10^{28} (p/m_p c)^{1/2} \text{ cm}^2 \text{ s}^{-1}$ and $\chi = 4$, with three sources producing CRs with kinetic energy $T = 1, 10$, and 100 GeV ; each source has luminosity $\mathcal{L} = 10^{38} \text{ erg s}^{-1}$; since the CRs from the different sources do not interact, the solution equation (104) applies independently to the CRs produced by each source. We run the simulation for $t = 2 \times 10^9 \text{ s}$ using a packet injection rate 10^{-3} s^{-1} , so that there are 2×10^6 CR packets at the end of the simulation. We show the results in Figs 2 and 3.

As the plots show, CRIPTIC recovers the exact solution to very high precision, covering ≈ 8 orders of magnitude in CR energy density; errors in the solution are consistent with expectations from Poisson statistics given the finite number of CR packets used in the simulation. The L^1 error, defined by

$$L_{\text{err}}^1 = \frac{4\pi}{L t} \int |U_{\text{CR,exact}} - U_{\text{CR,sim}}| r'^2 dr', \quad (105)$$




Figure 2. Energy density U_{CR} as a function of effective radius r' in the anisotropic diffusion test (Section 4.1.1). Solid lines show the exact solution computed from equation (104) for the three sources producing CRs of energy ($T = 1, 10$, and 100 GeV ; blue, orange, and green), while circles show the CRIPTIC result; for clarity we plot only every other bin. Error bars indicate the 90 per cent confidence interval derived from the number of CRs in the bin, assuming the CRs number counts are Poisson-distributed.




Figure 3. Projected energy density $\int U_{\text{CR}} dz$ in the anisotropic diffusion test (Section 4.1.1), where energy density is computed only for CRs with energies $T = 1 \text{ GeV}$. Colour shows the projected energy density from 10^{-1} – $10^2 \text{ eV cm}^{-3} \text{ pc}$, and white points show the positions of individual CR packets in regions where the CR energy density falls below this level.

where $U_{\text{CR,exact}}$ and $U_{\text{CR,sim}}$ correspond to the exact and simulation solutions shown in Fig. 2 (using the bins shown to compute $U_{\text{CR,sim}}$), is below 1 per cent for all three sources. Visual inspection of the projected CR distribution also demonstrates that it is anisotropic by a 4:1 ratio, exactly as expected.

4.1.2 Variable diffusion

Our next test investigates CRIPTIC's performance when the diffusion coefficient is non-constant. We simulate the transport of CRs with an isotropic diffusion coefficient, but one that varies as a power law in both space and time:

$$K = K_{\parallel} = K_{\perp} = K_0 \left(\frac{r}{r_0} \right)^{q_r} \left(\frac{t}{t_0} \right)^{q_t}. \quad (106)$$




Figure 4. Results of the variable diffusion coefficient test described in Section 4.1.2. We show the CR energy density U_{CR} in the initial condition (blue) at $t = t_0$ and in the final simulation snapshots at $t = 2t_0$ (orange). Solid lines show the exact solution given by equation (108), while circles show the CRIPTIC results, with error bars to indicate the 90 per cent confidence interval assuming the number of CR packets in each radial bin is Poisson-distributed.

There is no streaming in this test, and the background gas is at rest. No sources are present, and all CRs have momentum p_0 . For this setup the FPE is

$$\frac{\partial \tilde{f}}{\partial t} = K \left(\frac{\partial^2 \tilde{f}}{\partial x^2} + \frac{\partial^2 \tilde{f}}{\partial y^2} + \frac{\partial^2 \tilde{f}}{\partial z^2} \right), \quad (107)$$

and one can verify by substitution that the system admits a similarity solution for the spatial distribution

$$\tilde{f}(r, t) = N_{\text{CR}} \delta(p - p_0) \frac{3(2 - q_r) \frac{6}{q_r - 2}}{4\pi r_0^3 \Gamma\left(\frac{q_r - 5}{q_r - 2}\right)} [(q_t + 1)\eta]^{3/(2 - q_r)} \left(\frac{t}{t_0} \right)^{3 \frac{q_t + 1}{q_r - 2}} \exp \left[- \frac{(q_t + 1)\eta}{(q_r - 2)^2} \left(\frac{r}{r_0} \right)^{2 - q_r} \left(\frac{t}{t_0} \right)^{-q_t - 1} \right], \quad (108)$$

where $\Gamma(x)$ is the Γ function, $\eta = r_0^2/K_0 t_0$, and N_{CR} is the total number of CRs. The energy density $U_{\text{CR}} = T \int \tilde{f} dp$, where T is the kinetic energy corresponding to CR momentum p_0 .

For our test we set $K_0 = 4 \times 10^{28} \text{ cm}^2 \text{ s}^{-1}$, $r_0 = 10^{18} \text{ cm}$, $t_0 = 10^8 \text{ s}$, $q_r = -1/2$, and $q_t = 1$. We initialize the simulation with 2.5×10^5 CR packets with a total energy of $E_0 = N_{\text{CR}} T = 10^{48} \text{ erg}$; the initial radial distribution of the packets follow the exact solution, equation (108), evaluated at $t = t_0$. We then use CRIPTIC to advance the system to time $t = 2t_0$.

We show the results in Fig. 4. The figure shows that CRIPTIC recovers the exact solution to very high accuracy; errors are consistent with Poisson sampling, and are no larger in the final time step than in the initial setup. The L^1 error at the final time, defined by

$$L_{\text{err}}^1 = \frac{4\pi}{E_0} \int |U_{\text{CR,exact}} - U_{\text{CR,sim}}| r^2 dr, \quad (109)$$

is below 1 per cent.

4.1.3 Oscillating field loop

Our next test checks the ability of the code to trace diffusion along curved, moving field lines. In this test, we place a single CR source with CR luminosity $\mathcal{L} = 10^{38} \text{ s}^{-1}$ at $x = r_0 = 1 \text{ pc}$, $y = z = 0$ at time $t = 0$. The magnetic field consists of loops around the z -axis,




Figure 5. Results for the oscillating loop diffusion test (Section 4.1.3). The figure shows the CR energy density per unit angle $dE_{\text{CR}}/d\phi$, normalized by the total CR energy injected $\mathcal{L}t$. The solid line is the exact solution, and points show the numerical solution computed with CRIPTIC; as in previous tests, the points include error bars showing the 90 per cent confidence interval from Poisson statistics, but for this test the error bars are for the most part comparable to the sizes of the plot markers and thus are hidden.

so $\mathbf{B} = B_0 \hat{\phi}$, where $\hat{\phi}$ is the ϕ unit vector in an (r, ϕ, θ) cylindrical coordinate system. CRs diffuse with zero perpendicular diffusion, and parallel diffusion described by a coefficient $K_{\parallel} = 10^{28} \text{ cm}^2 \text{ s}^{-1}$. The FPE governing the system then reduces to

$$\frac{\partial \tilde{f}}{\partial t} = \frac{K_{\parallel}}{r_0^2} \frac{\partial^2 \tilde{f}}{\partial s^2} + \frac{\mathcal{L}}{T} \delta(s) \delta(p - p_0), \quad (110)$$

where $T = 1 \text{ GeV}$ is the energy of a single CR, p_0 is the corresponding CR momentum, and $s = 2\pi r_0 \phi$ is the position along the current loop, with the source located at $s = \phi = 0$.

The FPE in this case is equivalent to 1D diffusion in a periodic domain $(-\pi r_0, \pi r_0)$. This problem may be solved by standard Fourier methods, and the exact solution for a point source located at $\phi = 0$ that begins injecting CRs at $t = 0$ is that the CR energy per unit angle along the loop is

$$\frac{dE_{\text{CR}}}{d\phi} = \frac{\mathcal{L}t}{2\pi} \left[1 - 2 \frac{t_{\text{diff}}}{t} \sum_{n=1}^{\infty} \frac{\cos(n\phi)}{n^2} \left(e^{-n^2 t/t_{\text{diff}}} - 1 \right) \right], \quad (111)$$

where $t_{\text{diff}} = r_0^2/K_{\parallel} = 9.55 \times 10^8 \text{ s}$ is the characteristic diffusion time-scale. To add a complication to this test, both the source and the background gas perform simple harmonic oscillation in the x direction, with amplitude r_0 and angular frequency $\varpi = 4\pi/t_{\text{diff}}$; this allows us to test how well CRs follow field lines when the field lines are attached to a fluid that is accelerating, and where the direction of the acceleration can be both perpendicular and parallel to the field line. The exact solution is still given by equation (111), provided that we define ϕ relative to the time-dependent centre of the loop, since advection should move all the CR packets together.

We simulate the system for a time $t = 4t_{\text{diff}}$ using a packet injection rate $3 \times 10^{-4} \text{ s}^{-1}$, so over the course of the simulation 3.81×10^5 packets are injected into the domain, and the loop performs eight full periods of oscillation. This test therefore evaluates not only how well CRs follow curved field lines, but how well they do so when the field lines and the gas to which they are anchored are moving at arbitrary, time-variable angles relative to the field direction.

Fig. 5 shows the comparison between the CRIPTIC numerical solution and the exact solution given by equation (111). Clearly

the agreement is very good. Defining the L^1 error by

$$L^1_{\text{err}} = \frac{1}{\mathcal{L}t} \int \left| \frac{dE_{\text{CR,exact}}}{d\phi} - \frac{dE_{\text{CR,sim}}}{d\phi} \right| d\phi, \quad (112)$$

we find that the error is < 1 per cent. We also check quantitatively how well the CR packets remain confined to the field line to which they are attached, by examining the standard deviation of CR packet radial coordinates at the end of the simulation, σ_r . We find that $\sigma_r/r_0 \approx 5.7 \times 10^{-6}$, so over 4 diffusion times and 8 oscillations periods, numerical diffusion causes CR packets to drift across field lines by less than 1 part in 10^5 .

4.1.4 Momentum diffusion

Our next test evaluates the ability of the code to handle momentum diffusion, or, equivalently, second-order Fermi acceleration. For this test we turn on a single source of CRs at $t = 0$ at the origin, which injects CRs with luminosity $\mathcal{L} = 10^{38} \text{ erg s}^{-1}$. All injected CRs have a momentum of exactly $p_0 = 1 \text{ GeV}/c$. There is no spatial diffusion or streaming, but CRs diffuse in momentum with a diffusion coefficient K_{pp} , which we set implicitly by setting the diffusion time at momentum p_0 to $t_{\text{diff}} = 1 \text{ Myr}$. The corresponding diffusion coefficient is $K_{\text{pp}} = p_0^2/t_{\text{diff}}$, and the corresponding FPE is

$$\frac{\partial \tilde{f}}{\partial t} = K_{\text{pp}} \left[\frac{\partial^2 \tilde{f}}{\partial p^2} - \frac{\partial}{\partial p} \left(\frac{2}{p} \tilde{f} \right) \right] + \frac{\mathcal{L}}{T} \delta(\mathbf{r}) \delta(p - p_0), \quad (113)$$

where T is the kinetic energy corresponding to momentum p_0 . Using the change of variables $\tilde{f}' = p \tilde{f}$ reduces the problem to a 1D diffusion equation, subject to the boundary condition $\tilde{f}' = 0$ at $p = 0$, which can be solved by standard Green's function methods. The exact solution for the CR momentum distribution is

$$\begin{aligned} \frac{dn_{\text{CR}}}{dp} &= \frac{\mathcal{L}}{T} \frac{t}{\sigma_p} \frac{p}{p_0} \\ &\left[\sqrt{\frac{2}{\pi}} \left(e^{-\left(\frac{p-p_0}{2\sigma_p}\right)^2} - e^{-\left(\frac{p+p_0}{2\sigma_p}\right)^2} \right) + \frac{p+p_0}{\sigma_p} \text{erfc} \left(\frac{p+p_0}{\sqrt{2}\sigma_p} \right) + \right. \\ &\left. \frac{|p-p_0|}{\sigma_p} \text{erfc} \left(\frac{|p-p_0|}{\sqrt{2}\sigma_p} \right) - 1 \right], \end{aligned} \quad (114)$$

where $\sigma_p^2 = 2K_{\text{pp}}t$.

We simulate this problem with CRIPTIC using a packet injection rate of 10^{-7} s^{-1} , running to time $t = t_{\text{diff}} = 1 \text{ Myr}$, so there are 3.16×10^6 CR packets present at the end of the simulation. We compare the exact and numerical solutions in Fig. 6. The figure shows very good agreement between the exact and numerical results. Quantitatively, the L^1 error, defined for this problem as

$$L^1_{\text{err}} = \frac{1}{(\mathcal{L}/T)t} \int \left| \frac{dn_{\text{CR,exact}}}{dp} - \frac{dn_{\text{CR,sim}}}{dp} \right| dp, \quad (115)$$

is < 1 per cent.

4.1.5 Non-linear diffusion

All of our tests thus far have been for problems where the diffusion coefficient does not depend on the CR field, and thus which can be solved without the reconstruction step in our algorithm. We now consider a problem where reconstruction is required. We consider a system where diffusion is isotropic, and the diffusion coefficient is a power-law function of the CR energy density:

$$K_{\parallel} = K_{\perp} = K = K_0 \left(\frac{U_{\text{CR}}}{U_{\text{CR},0}} \right)^q. \quad (116)$$




Figure 6. Results for the momentum diffusion test (Section 4.1.4). The plot shows the CR momentum distribution dn_{CR}/dp at the end of the simulation. The solid line shows the exact solution, equation (114), while points show the CRIPTIC result. Error bars on the points show the 90 per cent confidence interval, computed assuming the number of CRs in each bin is Poisson-distributed.

No sources are present. Such a system admits a similarity solution (Pattel 1959)

$$U_{\text{CR}} = \frac{E_{\text{CR,tot}}}{\sqrt{\pi} r_0^3} \frac{\Gamma\left(\frac{1}{q} + \frac{5}{2}\right)}{\Gamma\left(\frac{1}{q} + 1\right)} \left[1 - \left(\frac{r}{r_{\text{out}}} \right)^2 \right] \left(\frac{t}{t_0} \right)^{-3/(3q+2)}, \quad (117)$$

where $E_{\text{CR,tot}}$ is the total CR energy in the system,

$$r_{\text{out}} = r_0 \left(\frac{t}{t_0} \right)^{1/(3q+2)}, \quad (118)$$

and r_0 and t_0 are given by

$$r_0^3 = \frac{E_{\text{CR,tot}}}{U_{\text{CR},0}} \pi^{-3/2} \frac{\Gamma\left(\frac{1}{q} + \frac{5}{2}\right)}{\Gamma\left(\frac{1}{q} + 1\right)} \quad (119)$$

$$t_0 = \frac{qr_0^2}{2(3q+2)K_0}. \quad (120)$$

Note that the system has the property that $U_{\text{CR}} = 0$ exactly at $r \geq r_{\text{out}}$, so this problem represents a severe test of our method, since the exact solution has a sharp cutoff in the CR energy density.

For our test with CRIPTIC, we set the total CR energy to $E_{\text{CR,tot}} = 10^{46} \text{ erg}$, adopt index $q = 1$, and specify $U_{\text{CR},0}$ and K_0 implicitly by setting $r_0 = 1 \text{ pc}$ and $t_0 = 10^7 \text{ s}$. We initialize the distribution of CR packets to the analytic solution at $t = t_0$, using a total of 10^5 CR packets, and evolve the system to $t = 3t_0$. We show the results of this test in Fig. 7. We find that the agreement between the numerical and exact solutions is very good. At smaller radii the match is almost perfect, and CRIPTIC recovers the location of the sharp edge of the exact solution with only a very small amount of numerical blurring. Defining the L^1 error for this test as

$$L^1_{\text{err}} = \frac{4\pi}{E_{\text{CR,tot}}} \int |U_{\text{CR,exact}} - U_{\text{CR,sim}}| r'^2 dr', \quad (121)$$

we find $L^1_{\text{err}} = 2.4$ per cent.

To understand the quality of the solution, it is helpful to examine the reconstructed CR energy density and its derivative, which are used to construct the diffusion coefficient and its spatial derivative. We show these reconstructions for the initial time in Fig. 8; results are qualitatively similar at later times. We see that our kernel density




Figure 7. Results for the non-linear diffusion test (Section 4.1.5). The plot shows the CR energy density U_{CR} at times $t = t_0$ ('Initial', blue) and $t = 3t_0$ ('Final', orange). Lines show the exact solution given by equation (117), while circles with error bars show the solution computed by CRIPATIC; errors indicate the 90 per cent confidence interval derived by assuming the number of CR packets in each bin is drawn from a Poisson distribution.




Figure 8. Reconstruction of the CR energy density U_{CR} (top) and its radial derivative dU_{CR}/dr (bottom) at the initial time in the non-linear diffusion test (Section 4.1.5). Black lines show the exact solution as a function of radius r , solid blue lines show the median value of the CRIPATIC CR packets, computed in 30 radial bins, and blue and orange bands show the 25th–75th percentile range (50 per cent confidence interval) and 5th–95th percentile range (90 per cent CI), respectively, in the same bins.

estimation algorithm reconstructs the CR energy density with very high accuracy; the median is so close to the exact solution that the line showing it is nearly invisible in the figure, hidden behind the exact solution, and the 50th and 90th percentile ranges lie within ≈ 10 per cent and ≈ 20 per cent of the exact solution except at large radii.

The derivative, which is inherently harder to reconstruct, shows a larger range, and deviates from the exact solution at both small radii, where the number of CR packets is small due to the small volume, and at the edge of the distribution; none the less, the median agrees very well with the exact value over most of the radial range. Errors

are most significant at the largest radii, where the exact solution goes to zero exactly, a feature that is necessarily blurred somewhat in CRIPATIC's reconstruction due to the finite size of the kernel. It is this effect that is responsible for the sharp edge of the exact solution expanding very slightly too quickly in the CRIPATIC simulation. None the less, the overall error is very small.

4.1.6 Streaming and streaming losses

Our next test checks the ability of the code to handle streaming down CR pressure gradients, together with the associated adiabatic changes in CR momentum when the divergence of the streaming velocity is non-zero. For this test we consider a fully ionized medium at rest with a power-law distribution of density and a 'split monopole' magnetic field. The density and magnetic field as a function of position are

$$\rho = \rho_0 \left(\frac{r}{r_0} \right)^{k_\rho} \quad (122)$$

$$\mathbf{B} = B_0 \left(\frac{r}{r_0} \right)^2 \text{sgn}(z) \hat{\mathbf{r}}, \quad (123)$$

where r is the distance from the origin. For this test we use $r_0 = 1$ pc, $\rho_0 = 2.34 \times 10^{-24}$ g cm $^{-3}$, $k_\rho = -2$, and $B_0 = 10$ μ G. The simulation begins with no CRs, but we place a point source of CRs with luminosity $\mathcal{L} = 10^{38}$ erg s $^{-1}$ at the origin, where it injects CRs with initial energy $T = 1$ GeV.¹⁰ For the CR transport model in this set, we let CRs stream down the CR pressure gradient at the ion Alfvén velocity,

$$\mathbf{w} = \frac{\mathbf{B}}{\sqrt{4\pi\rho}} \text{sgn}(-\nabla P_{\text{CR}} \cdot \mathbf{B}). \quad (124)$$

In this test the CR pressure gradient always points to the origin, and thus CRs should always stream away from the origin at the streaming speed

$$w = v_{A0} \left(\frac{r}{r_0} \right)^{-2-k_\rho/2}, \quad (125)$$

where $v_{A0} = B_0/\sqrt{4\pi\rho_0}$. Individual CRs therefore obey an equation of motion $dr/dt = w$ which, for a CR injected at $t = t_{\text{inj}}$, has the solution

$$r(t - t_{\text{inj}}) = r_0 \left[\left(3 + \frac{k_\rho}{2} \right) \frac{v_{A0} t}{r_0} \right]^{1/(3+k_\rho/2)} \quad (126)$$

for times $t \gg t_{\text{inj}}$. Moreover, the momentum of individual CRs evolves with radius due to adiabatic cooling as

$$\frac{dp}{dr} = \left(\frac{dr}{dt} \right)^{-1} \frac{dp}{dt} = -\frac{p}{3w} \nabla \cdot \mathbf{w}, \quad (127)$$

which has the solution that a CR injected with momentum p_{inj} at radius r_{inj} has momentum

$$p = p_{\text{inj}} \left(\frac{r}{r_{\text{inj}}} \right)^{-k_\rho/6} \quad (128)$$

once it reaches radius r .

¹⁰To avoid CRs near the origin requiring infinitely small time steps, we flatten the density and magnetic field profiles at very small radii, and make the CR source slightly extended. Specifically, we adopt a flattening radius $r_{\text{flat}} = 10^{-3} r_0$ and evaluate the density and magnetic field using $\max(r, r_{\text{flat}})$ rather than simply r ; we likewise inject CRs at a random radius uniformly distributed from 0 to r_{flat} , rather than exactly at $r = 0$.




Figure 9. Solutions for the streaming problem (Section 4.1.6). In the top panel, we show CR packet radial position r versus packet age $t_{\text{packet}} = t - t_{\text{inj}}$, and in the bottom we show CR packet momentum p as a function of radius. In both panels, solid lines are the exact solutions given by equations (126) and (128), respectively, and circles with error bars are the numerical solutions computed by CRIPTIC, averaged over bins in age (top) and radius (bottom). Error bars show the 5th to 95th percentile range in each bin.

We simulate the system for 10^9 s using a packet injection rate of 10^{-4} s^{-1} , so that by the end of the simulation there are 10^5 packets. We show the CRIPTIC results in comparison to the exact solutions for CR packet position versus age and momentum versus radius in Fig. 9. We again see excellent agreement between CRIPTIC and the exact solution, indicating the code successfully reconstructs the direction of the CR pressure gradient and correctly computes the rate of adiabatic cooling caused by the non-uniform streaming speed. Quantitatively, we define the radial position and momentum error for each CR packet by

$$e_r = \frac{r}{r_{\text{exact}}} - 1 \quad e_p = \frac{p}{p_{\text{exact}}} - 1, \quad (129)$$

where r_{exact} and p_{exact} are the exact solutions given by equations (126) and (128). We find that the mean values of e_r and e_p are 0.37 and 0.40 per cent, respectively, with variances 1.5 and 4.6 per cent; the errors are therefore small, and in fact the variance in e_p is an overestimate of the true error because it is mostly a result of the dependence of p on the initial injection radius r_{inj} , which is randomly varied by a small amount for numerical reasons as discussed above.




Figure 10. Number density of CRs with momentum $> p$ in the reconstruction test described in Section 4.1.7. Solid lines show the median value of $n_{\text{CR}}(> p)$ computed by CRIPTIC in 25 logarithmically-spaced bins of p , evaluated at 5 different radii (as indicated by the colours). Shaded bands around the medians show the ranges within which 50 and 90 per cent of the reconstructed values fall. Dashed lines show the exact solution (equation 130).

4.1.7 Number density computation

Our final transport test evaluates CRIPTIC's ability to reconstruct the CR field in a more realistic problem where there is a continuous distribution of CR positions and momenta. In this problem we place a source of CR protons at the origin, which injects CRs at a specific rate $d\dot{n}_{\text{src}}/dp \propto p^q$ over a momentum range from p_0 to p_1 . The CRs then diffuse isotropically away from the source with a constant diffusion coefficient K . The exact solution for the CR distribution is then just given by equation (104) with the $(\mathcal{L}/T)\delta(p - p_0)$ replaced by $d\dot{n}_{\text{src}}/dp$, and we can immediately write down the expected number density of CRs at any given radius r and time t with momentum $> p$ for any $p \in (p_0, p_1)$,

$$n_{\text{CR}}(> p) = \frac{1}{4\pi K r} \text{erfc}\left(\frac{r}{2\sqrt{Kt}}\right) \dot{n}_{\text{src}} \frac{1 - (p/p_1)^{q+1}}{1 - (p_0/p_1)^{q+1}}, \quad (130)$$

where $\dot{n}_{\text{src}}$ is the total CR injection rate integrated over all momenta. The test is to compare this exact value for the CR number density to the value reconstructed by CRIPTIC, given by equation (74). We run the test using a diffusion coefficient $K = 10^{28} \text{ cm}^2 \text{ s}^{-1}$ and a source with total energy injection rate $\mathcal{L} = 10^{38} \text{ erg s}^{-1}$, with $p_0 = 1 \text{ GeV}/c$, $p_1 = 10^3 \text{ GeV}/c$, and $q = -2.2$, for a time, $t = 10^9 \text{ s}$. We use a packet injection rate 10^{-4} s^{-1} for the test.

We compare the exact and reconstructed CR number densities in Fig. 10, which shows $n_{\text{CR}}(> p)$ computed in 25 logarithmically spaced bins in p , evaluated at radii $r/r_{\text{diff}} = 0.2, 0.5, 1.0, 2.0$, and 3.5 , where $r_{\text{diff}} = \sqrt{Kt}$, and for each radial bin we consider CR packets whose radii are within 5 per cent of the target value. Solid lines show the median reconstructed value of $n_{\text{CR}}(> p)$ for packets in that bin, and shaded bands show the ranges within which 50 and 90 per cent of the packet values fall. We see that the median reconstructed value of $n_{\text{CR}}(> p)$ is in almost perfect agreement with the exact result except at the highest values of p , where the finite number of sample packets causes deviation, and at $r/r_{\text{diff}} = 0.2$, where the simulation median is ~ 10 per cent below the true value due to the finite size of the smoothing kernel, which blurs out the sharp peak at small radii visible in, e.g. Fig. 2. The other clear trend is that the 50 and 90 per cent ranges expand at larger radii, simply because our finite number of packets leads to larger Poisson errors.




Figure 11. Distribution of exact (horizontal axis) and CRIPTIC-computed (vertical axis) values of $n_{\text{CR}}(>p)$, the number density of CRs with momentum $>p$. The blue solid line shows the median, while shaded regions indicate the range into which 50 and 90 per cent of the reconstructed values fall; the black dashed line is the 1:1 line, corresponding to perfect agreement.

To make a quantitative analysis of the error, for each CR packet we evaluate the exact value of $n_{\text{CR, exact}}(>p)$ at its position from equation (130), and we place the packets in 40 bins of $n_{\text{CR, exact}}(>p)$. For each bin, we examine the distribution of CRIPTIC-reconstructed number densities, $n_{\text{CR, sim}}(>p)$, in that bin, and compute the median and the 50 and 90 per cent confidence intervals of this distribution. We plot the median and confidence intervals as a function of $n_{\text{CR, exact}}(>p)$ in Fig. 11. We see the same patterns that were visible in Fig. 10, i.e. agreement is excellent over most of the range of $n_{\text{CR, exact}}(>p)$, but there are deviations at both the highest values, where the smoothing kernel blurs out the very sharp peak around the source, and the lowest values, where the finite number of CR packets in the simulation leads to errors in reconstructing very low density regions. However, over 5 decades in $n_{\text{CR}}(>p)$, from 10^{-13} to 10^{-8} cm^{-3} , the error is very small: averaged over this range, the median value differs from the exact one by only 0.005 dex, and the 50 per cent range is only 0.16 dex wide.

4.2 Microphysics tests

Our second set of tests evaluates the performance of our code in simulating the microphysical processes that govern CR loss and secondary production, and the radiation spectra produced thereby. In only some of these cases is an exact analytical solution available, and, where it is not, we compare to expected physical behaviour and limiting cases.

4.2.1 Proton diffusion with collisional loss

Our first test of microphysics evaluates our treatment of proton catastrophic losses, and the coupling between them and transport. To this end, we repeat the anisotropic diffusion experiment described in Section 4.1.1, but including losses due to CR proton inelastic scattering due to a background gas of constant density ρ , and using a source that injects a continuous momentum distribution of CR protons at a rate per unit momentum $d\dot{n}_{\text{src}}/dp$; we disable all other loss mechanisms for the purposes of this test. Considering only primary protons (i.e. those that have not yet been scattered), the

FPE that governs this system is

$$\frac{\partial \tilde{f}}{\partial t} = K \left(\chi \frac{\partial^2 \tilde{f}}{\partial x^2} + \frac{\partial^2 \tilde{f}}{\partial y^2} + \frac{\partial^2 \tilde{f}}{\partial z^2} \right) + \frac{d\dot{n}_{\text{src}}}{dp} \delta(\mathbf{r}) - \sigma_{\text{nuc}} v \frac{\rho}{\mu_{\text{H}} m_{\text{H}}}, \quad (131)$$

where K is the perpendicular diffusion coefficient, χ is the ratio of parallel and perpendicular coefficients (with the magnetic field oriented in the $\hat{x}$ direction), and σ_{nuc} and v are the nuclear inelastic cross-section and particle velocity as a function of particle momentum p . The system can be solved exactly by making the same change of variable $x = \sqrt{\chi} x'$ and $r'^2 = x'^2 + y^2 + z^2$ as in Section 4.1.1 to transform the problem to a standard diffusion equation with a loss term, and then writing down the Green's function including the loss term. Since the loss rate is independent of position and time, this is simply

$$G(r', t) = \frac{1}{(4\pi K t)^{3/2}} \exp \left(-\frac{r'^2}{4K t} - \frac{t}{t_{\text{loss}}} \right), \quad (132)$$

where $t_{\text{loss}} = \mu_{\text{H}} m_{\text{H}} / \sigma_{\text{nuc}} v \rho$ is the loss time-scale. The exact solution (again, considering only the primary proton population), is

$$\begin{aligned} \tilde{f}(r', p, t) &= \int_0^t \sqrt{\chi} \frac{d\dot{n}_{\text{src}}}{dp} G(r', t') dt' \\ &= \frac{d\dot{n}_{\text{src}}}{dp} \frac{1}{8\pi \sqrt{\chi} K r'} \left[e^{-r'/r_{\text{loss}}} \text{erfc} \left(\frac{r'}{2r_{\text{diff}}} - \frac{r_{\text{diff}}}{r_{\text{loss}}} \right) + \right. \\ &\quad \left. e^{r'/r_{\text{loss}}} \text{erfc} \left(\frac{r'}{2r_{\text{diff}}} + \frac{r_{\text{diff}}}{r_{\text{loss}}} \right) \right], \end{aligned} \quad (133)$$

where we have defined $r_{\text{loss}}^2 = K t_{\text{loss}}$ and $r_{\text{diff}}^2 = K t$.

We run the test using a momentum-dependent diffusion coefficient $K = 10^{28} (p/m_p c)^{1/2} \text{ cm}^2 \text{ s}^{-1}$ and anisotropy parameter $\chi = 4$. The central CR source produces CR protons with a momentum distribution $d\dot{n}_{\text{src}}/dp \propto p^q$ with $q = -2.2$ over a momentum range from $p_0 = 0.1 \text{ GeV}/c$ to $p_1 = 10^5 \text{ GeV}/c$, with a total luminosity $L = 10^{38} \text{ erg s}^{-1}$. The CRs propagate through a uniform background gas of density $\rho = 2.34 \times 10^{-21} \text{ g cm}^{-3}$, and we run the simulation for $t = 2 \times 10^{12} \text{ s}$, using a packet injection rate $\Gamma = 10^{-6} \text{ s}^{-1}$, so there are 2×10^6 packets present at the end of the simulation. Given this setup, we have $r_{\text{diff}} = 45.8 (p/m_p c)^{1/2} \text{ pc}$, and $r_{\text{loss}}/r_{\text{diff}} \approx 0.5$ at momenta far above the pion production threshold, $p_{\text{th}} = 0.78 \text{ GeV}/c$; thus we have selected parameters so that losses are relatively important over most of the momentum range of the test, but become unimportant ($r_{\text{loss}} \rightarrow \infty$) at the lowest CR momenta.

We show the radial distribution of CR number density for a range of sample momenta in Fig. 12, comparing the simulation results to the exact solution given by equation (133); for comparison we also show the solution that would be expected in the absence of losses, i.e. setting $r_{\text{loss}} = \infty$ in equation (133). We find that CRIPTIC recovers the correct exact solution, including momentum-dependent loss rates, to the level expected from Poisson statistics. We define the L^1 error norm for this problem by

$$L_{\text{err}}^1 = \frac{\int_0^\infty \int_{p_0}^{p_1} |\tilde{f}_{\text{sim}} - \tilde{f}_{\text{exact}}| 4\pi r'^2 p^q dp dr'}{t \int_{p_0}^{p_1} \frac{d\dot{n}_{\text{src}}}{dp} p^q dp}, \quad (134)$$

where $\tilde{f}_{\text{sim}}$ are the simulation results and $\tilde{f}_{\text{exact}}$ is the exact solution. Note the weight factor of p^q in this integral is included to ensure that all momenta are weighted equally in computing the error estimate, so the integrand in the denominator is independent of p . Defined this way, we find that the L^1 error in our CRIPTIC solution is 2.0 per cent.




Figure 12. CR proton number density per unit momentum, dn_{CR}/dp , as a function of effective radius r' in the anisotropic diffusion with loss test problem (Section 4.2.1). We show dn_{CR}/dp evaluated in five sample momentum bins, each 0.2 dex wide, spanning from the minimum to the maximum momenta present in the problem; note that our plotted values of dn_{CR}/dp have been scaled by p^3 in order to make it easier to display such a wide range of momenta on the same plot. Points with error bars show CRIPTIC simulation results, with the error bar indicating the Poisson uncertainty on the mean value in each bin due to the finite number of packets. Solid lines show the exact solutions given by equation (133), while dashed lines show the solutions we would expect to find if we were to disable losses (equation 133 with $n_{\text{loss}} \rightarrow \infty$).

4.2.2 Electron streaming with synchrotron and inverse Compton loss

Our second test evaluates our implementation of synchrotron and inverse Compton losses for electrons. We place a point source of CR electrons with a powerlaw momentum distribution $d\dot{n}_{\text{src}}/dp = (\dot{n}_0/p_0)(p/m_e c)^q$ over the range (p_0, p_1) in a medium with a uniform magnetic field $\mathbf{B} = B\hat{x}$ and a uniform radiation field with dilution factor W_{BB} and blackbody temperature T_{BB} . The electrons stream in the $+x$ direction at constant speed w , while suffering synchrotron and inverse Compton losses. All other loss processes are disabled, as is diffusion. The Fokker–Planck equation for this system is

$$\frac{\partial \tilde{f}}{\partial t} = -\frac{\partial}{\partial x}(w\tilde{f}) + \frac{\partial}{\partial p}(\dot{p}_{\text{cts}}\tilde{f}) + \frac{d\dot{n}_{\text{src}}}{dp}, \quad (135)$$

where $\dot{p}_{\text{cts}}$ is the rate of momentum loss due to synchrotron and inverse Compton radiation. Since the loss rate is independent of position, the spatial and momentum parts of the system are separable, and there is a one-to-one relationship between the position x and the time $t = x/w$ for which CRs have been subject to loss processes. Over this time, a CR injected with momentum p_i will have been reduced to momentum p given implicitly by $t = -\int_{p_i}^p dp'/\dot{p}_{\text{cts}}(p')$. If we further adopt the ultrarelativistic limit $\gamma \gg 1$ and $p \approx \gamma m_e c$ and assume that we are far from the Klein–Nishina regime, then $\dot{p}_{\text{cts}} \propto p^2$, and we can evaluate the integral analytically; from equation (39) and equation (59), we have

$$\gamma = \frac{\gamma_i}{1 + \gamma_i(t/t_{\text{loss}})}, \quad (136)$$

where for convenience we have expressed the momentum in terms of the Lorentz factor γ , and

$$t_{\text{loss}} = \frac{3m_e c}{4\sigma_T (U_B + U_R)}, \quad (137)$$

where U_B and U_R are the energy densities of the magnetic field and radiation field, respectively. This in turn allows us to write down the solution to the Fokker–Planck equation,

$$\tilde{f}(x, \gamma, t) = \frac{\dot{n}_0}{\gamma_0} \left(\frac{1}{1 - \xi\gamma} \right)^2 \left(\frac{\gamma}{1 - \xi\gamma} \right)^q \quad (138)$$

for $\xi < t/t_{\text{loss}}$ and $\gamma \in [\gamma_0/(1 + \xi\gamma_0), \gamma_1/(1 + \xi\gamma_1)]$, and 0 otherwise, where here $\gamma = p/m_e c$ is the Lorentz factor in the ultrarelativistic limit and $\xi \equiv x/wt_{\text{loss}}$ is the dimensionless distance. Integrating over x , the momentum distribution of all electrons with $\gamma > \gamma_0$ is

$$\frac{dn_{\text{CR}}}{d\gamma} = \frac{\dot{n}_0 t_{\text{loss}}}{\gamma_0} \left(\frac{\gamma^{q-1}}{q+1} \right) [(1 - \xi_{\text{max}}\gamma)^{-q-1} - 1], \quad (139)$$

where $\xi_{\text{max}} = \min(t/t_{\text{loss}}, \gamma^{-1} - \gamma_1^{-1})$; for $1 \ll \gamma \ll \gamma_1$ and $t \rightarrow \infty$, this gives the classical ‘cooled synchrotron’ $dn/dp \propto p^{q-1}$, i.e. the spectral index is the injection index minus one.

We carry out this test with a source with total luminosity 10^{38} erg s $^{-1}$, which injects CR electrons with a power-law momentum distribution characterized by $q = -2.2$ at momenta from $p_0 = 10^{-2}$ GeV/c to $p_1 = 10^3$ GeV/c. The electrons stream at $w = 100$ km s $^{-1}$, and we set the energy densities in the background magnetic field and radiation field to $U_B = U_R = 50$ eV cm $^{-3}$; this corresponds to a magnetic field $B = 44.87$ μ G, a blackbody radiation temperature $T_{\text{BB}} = 10.144$ K (for dilution factor $W_{\text{BB}} = 1$), and a loss time $t_{\text{loss}} = 6.09$ Gyr. We run the simulation for $10^{-3}t_{\text{loss}}$, using a packet injection rate 10^{-8} s $^{-1}$. We show the results in Fig. 13; the upper panel shows the CR spectrum integrated over all positions compared to the analytical solution given by equation (139), while the lower panel shows the distribution function evaluated at selected positions, compared to the analytical solution given by equation (138).

As Fig. 13 shows, CRIPTIC reproduces the exact solutions very well. Quantitatively, we define the L^1 error for the integrated spectrum for this problem by

$$L^1_{\text{err}} = \frac{1}{\ln(\gamma_1/\gamma_0)} \int_{\ln \gamma_0}^{\ln \gamma_1} \left| \log_{10} \left(\frac{(dn_{\text{CR}}/d\gamma)_{\text{sim}}}{(dn_{\text{CR}}/d\gamma)_{\text{exact}}} \right) \right| d \ln \gamma, \quad (140)$$

so the error is the mean logarithmic deviation between the exact and simulated spectra; note that we use this definition because the steep nature of the spectrum means that, if we do not measure the deviation logarithmically, the error norm is dominated by the parts of the spectrum at low γ , where we simply have the original spectrum. Using this definition, we find $L^1_{\text{err}} = 0.012$ dex for the solution shown in Fig. 13. We can also define the L^1 error at a given dimensionless position ξ analogously, simply by replacing $dn_{\text{CR}}/d\gamma$ by $\tilde{f}$, and replacing γ_1 with $\gamma_1/(1 + \xi\gamma_1)$ in the upper integration limit; doing so we find $L^1_{\text{err}} = 0.026$ dex at $\xi = 10^{-5}$ to 0.01 dex at $\xi = 10^{-3}$; the error is largest at 10^{-5} because this is sampled by the fewest packets. Nonetheless, even at this small value of ξ , the agreement with the exact solution is clearly very good.

Although we do not have exact solution for it, we can also slightly modify this test to verify that CRIPTIC behaves as expected qualitatively in the Klein–Nishina regime. To do so, we change the background magnetic field to $B = 4.487 \times 10^{-9}$ μ G and the radiation field to $T_{\text{BB}} = 1.0144 \times 10^5$ K with $W_{\text{BB}} = 2 \times 10^{-16}$; this has the effect of setting $U_R = 100$ and $U_B = 5 \times 10^{-19}$ eV cm $^{-3} \approx 0$. Thus the total magnetic plus radiation energy density is unchanged from the original version of the test, and in the Thomson limit we should recover exactly the same solution. However, while for our previous value of T_{BB} the momentum range (p_0, p_1) for the injected CRs corresponded to $\log \Gamma_{\text{BB}} = -6.9$ to -1.9 , with the higher value of T_{BB} we now have $\log \Gamma_{\text{BB}} = -2.9$ to 2.1 . Consequently $\Gamma_{\text{BB}} = 1$ occurs in the middle of the injected momentum range, at $p = 7.5$




Figure 13. Results of the electron streaming with losses test. In the top panel, we show the compensated CR spectrum integrated over all space, $\gamma^2(dn_{CR}/d\gamma)$ at the end of the simulation. Filled circles show the numerical results evaluated in 30 logarithmically spaced bins of Lorentz factor γ , the solid line shows the exact solution (equation 139), the dashed grey line shows the distribution of CRs injected by the source, and the open circles are the results of our test in the Klein–Nishina regime (see main text). In the lower panel, we show momentum distributions evaluated at five sample positions given by the dimensionless position variable ξ indicated in the legend. As in the top panel, filled circles are the CRPTIC numerical solution, solid lines are exact solutions (equation 138), the grey dashed line is the injected spectrum, and the open circles are the results of running the simulation in the Klein–Nishina regime rather than the Thomson regime.

GeV/c ($\gamma = 1.5 \times 10^4$), and we therefore expect Klein–Nishina effects to become significant at momenta approaching this value.

The open circles in Fig. 13 show the results of the test in the Klein–Nishina regime. The qualitative result is as expected: for $\gamma \ll 1.5 \times 10^4$, we are in the limit $\Gamma_{BB} \ll 1$, and the solution matches the Thomson case. For $\gamma \gg 1.5 \times 10^4$, on the other hand, the rate of energy loss scales as $dE/dt \propto \ln \gamma$ (Blumenthal & Gould 1970), so the loss time-scale obeys $t_{loss} \propto \gamma/\ln \gamma$, rather than $t_{loss} \propto 1/\gamma$ as in the Thomson regime. Thus losses become increasingly unimportant at high γ , and the spectrum approaches the injected spectrum rather than the Thomson limit solution.

4.2.3 A thick target

Our final microphysical test is to simulate a thick target with all microphysical processes enabled, and with a source injecting both primary protons and electrons. In this test CRPTIC is performing a calculation similar to that carried out by other authors who treat CR microphysics but do not include transport, or include it only in a simplified parametrized way such as by analytically specifying a loss time or calorimetry fraction (e.g. Yoast-Hull et al. 2014; Peretti et al. 2019; Roth et al. 2021). Our goal is to show that CRPTIC, though not optimized for this type of calculation (since it is possible to obtain the answer much more efficiently if one is uninterested in




Figure 14. Steady-state CR spectra in the thick target problem. For each species we plot $T^2 dn_{CR}/dT$, where dn_{CR}/dT is the number of individual CR particles in a particular energy bin. Different colours indicate primary protons, primary electrons, secondary electrons, and (secondary) positrons, as indicated in the legend. The dashed black line labelled $T^{-2.2}$ shows the shape of the injection spectrum.

a detailed treatment of spatial transport), none the less recovers CR and emitted γ -ray spectra similar to those that have been reported in the literature.

For this test we disable spatial transport and momentum diffusion, and we consider an environment such as might be found in a starburst galaxy: a uniform medium of molecular hydrogen characterized by number density $n_H = 10^3$ H nuclei cm^{-3} , an ionization fraction by mass $\chi = 10^{-4}$, a magnetic field strength $B = 0.3$ mG, and two radiation fields both with $W_{BB} = 1$, one with $T_{BB} = 2.73$ K (the CMB) and one with $T_{BB} = 20$ K (representing a reprocessed dust radiation field). We place two sources in the medium. One injects protons with a momentum distribution $dn/dp \propto p^{-2.2}$ over a range in kinetic energy $T = [10^{-3}, 10^6]$ GeV, and has total luminosity $L = 3 \times 10^{42}$ erg s^{-1} , corresponding roughly to the luminosity expected for a galaxy with a star formation rate of $\approx 100 M_\odot \text{yr}^{-1}$, assuming 1 SN per $100 M_\odot$ of stars formed, a total energy of 10^{51} erg per SN, with ≈ 10 per cent of that taking the form of CRs. The other source injects electrons with the same spectrum, but a total luminosity a factor of 10 smaller. We run the simulation for 0.5 Myr using a primary packet injection rate $\Gamma_{inj} = 2 \times 10^{-7} \text{s}^{-1}$; this is longer than the loss time at all energies for the background environment, so by this time the system settles to steady state. At the end of the simulation, there are approximately 4.3×10^6 CR packets present.

We show the resulting steady-state CR spectra in Fig. 14. The result is in accord with what we would expect: the proton spectrum has a slope that is very slightly shallower than the injection spectrum at high energies, reflecting the slight increase in pp cross-section with energy. This continues to ≈ 1 GeV, and below this energy the spectrum dies off quickly, reflecting the strong ionization losses that low-energy CRs suffer in a starburst environment (c.f. Fig. 1); indeed, each second the CR protons ionize a total of 6.0×10^{50} H_2 molecules and 7.2×10^{49} He atoms.¹¹ The primary electron spectrum is both lower in absolute value and substantially steeper, with a slope closer to -3.2 at high energy, as a result of the quadratic dependence of

¹¹Note that these should be understood as primary ionizations, since we do not track electrons with energies $\lesssim 1$ MeV, which overwhelmingly dominate secondary ionizations.




Figure 15. Emitted γ -ray spectra produced in the thick target test. The black solid line shows the total spectrum produced by all particles and processes, and the blue solid line shows the contribution from nuclear inelastic scattering. Dashed and dot-dashed lines in different colours show emission by primary and secondary electrons and positrons, respectively, with the colour indicating the emission process – bremsstrahlung, inverse Compton, and synchrotron.

synchrotron and inverse Compton losses on CR energy. It too falls off below ≈ 1 GeV due to ionization losses – the total ionization rate produced by electrons and positrons is 6.8×10^{50} and $8.7 \times 10^{49} \text{ s}^{-1}$ for H_2 and He , respectively. Finally, the secondary electron and positron spectra are nearly identical, and are lower still, reflecting the relatively large value we have chosen for the primary electron/proton ratio in this test.

We show the γ -ray emission that CRIPTIC predicts for this system in Fig. 15. We find that nuclear inelastic scattering dominates at high energies, with a sharp cutoff just below 10^6 GeV reflecting the cutoff in the injection spectrum that we apply there. This gives rise to the usual bump at ≈ 1 GeV in the spectral energy distribution. Bremsstrahlung and inverse Compton emission from primary electrons significantly contribute to the total emission at energies below ≈ 1 GeV. In the radio, we see a dominant contribution from primary electron synchrotron, with a ≈ 10 per cent additional contribution from secondary electrons. All other processes are subdominant. Note that the wavy structure in the inverse Compton and bremsstrahlung spectra at high energy is real, and reflects the modulation imposed by the fact that we have two blackbody radiation fields at different temperatures present. On the other hand, the somewhat spiky contribution from positron annihilation (which includes only the Doppler-boosted 511 keV photons, not the inverse Compton, bremsstrahlung, or synchrotron contributions from positrons) is a result of the relatively small number of positron packets present in the calculation, which causes some numerical noise. However, since this component is subdominant by ≈ 2 dex, this has no noticeable effect on the total spectrum. In general, our result is qualitatively consistent with the results of other thick target calculations of starburst γ -ray and radio spectra.

5 DISCUSSION AND CONCLUSION

We conclude by discussing applications of CRIPTIC, as well as the limitations of the current code and our plans for future expansions.

5.1 Applications

CRIPTIC can be used to solve a wide range of problems in CR transport. In Sampson et al. (2022), we have already applied it to the problem of determining an effective transport theory for CRs that stream through a turbulent plasma. This application exploits CRIPTIC’s ability to model transport through an arbitrary, time-dependent background – in this case the output of an MHD turbulence simulation.

In future work we intend to use CRIPTIC to post-process MHD simulations of Milky Way-like galactic discs (e.g. Wibking & Krumholz 2022), in order to compare the results produced by different CR transport models with observable quantities such as the γ -ray spectral index as a function of height above the galactic plane. We also intend to post-process CR hydrodynamics simulations in order to predict detailed observables from them.

Astrochemistry and the link between it and high-energy phenomena represents another immediate application. There has been considerable debate about the CR ionization rate in starburst galaxies (e.g. Papadopoulos 2010; Bisbas, Papadopoulos & Viti 2015; Narayanan & Krumholz 2017) and in the Milky Way Central Molecular Zone (Ginsburg et al. 2016; Oka et al. 2019; Tanaka, Nagai & Kamegai 2021), but thus far this discussion has focused on molecular indicators, and has not taken advantage of constraints offered by high-energy tracers such as γ -ray emission. However, the CRs that drive ionization and those that produce γ -rays are ultimately part of the same population, albeit at somewhat different energies. Efforts to combine constraints on the low-energy population that drives ionization and the high-energy population that drives γ -ray emission have thus far been very limited (e.g. Armillotta et al. 2022), but CRIPTIC’s ability to simulate the full CR energy range should greatly facilitate these efforts.

5.2 Limitations and expansion plans

No simulation code can ever capture all of physics, and that is certainly true of CRIPTIC. It is therefore worth concluding pointing

out some limitations of the current code, and some plans for future expansion to remedy at least some of these limitations. One limitation is that CRIPTIC is a post-processing code, and therefore does not allow for calculation of the backreaction on the flow due to CR pressure forces or heating. In this regard, it is analogous to Monte Carlo post-processing radiative transfer codes, which can be used to predict detailed spectra, but do not allow self-consistent calculation of how a flow reacts to radiation forces. To the extent that a calculation using criptic predicts a CR field where CR forces are important (or, analogously, a Monte Carlo calculation predicts a dynamically important radiation pressure force), the correct way to proceed is to evolve the system with a code that self-consistently includes those forces, and then use a post-processing tool like CRIPTIC to calculate the observable emission at much higher resolution than would be possible from the self-consistent calculation alone. Beyond this limit to the overall CRIPTIC methodology, though, we here identify three other limitations that are ripe for improvement in future releases.

First, CRIPTIC does not yet include all of the CR species or loss processes for which observational constraints exist. Since we have focused on radiative and astrochemical signatures from galaxies, and particularly radiative signatures that are observable from beyond the Milky way, in this first release we have included the species that dominate these. However, direct *in situ* measurements exist for a range of heavier CR nuclei, most prominently He, B, and C. We intend to include these species in a future release. We will also add antiprotons and the process of positron annihilation via positronium formation, which do produce radiative signatures that are observable in the Galaxy, if not from external galaxies. Finally, we have focused on CRs in the ~ 1 MeV – 1 PeV range in typical galactic environments, and have not included loss processes that become dominant outside this range or in highly magnetized environments such as around active galactic nuclei. At low energies the main omitted process is charge exchange (e.g. Schultz et al. 2008), while at the high energy end it is photohadronic interactions (e.g. Mücke et al. 1999); photon–photon scattering can also become a significant loss process for high-energy photons in strong radiation environments, and is not yet included. In highly magnetized environment, we should also include synchrotron losses for protons. The modular nature of the code makes addition of such processes straightforward, and these too may be included in future releases.

Secondly, at present CRIPTIC solves the Fokker–Planck equation in the spatial and momentum directions, but not in pitch angle; it is therefore valid only on scales large enough that the local CR pitch angle distribution has become approximately isotropic. This limits applications to ultrahigh energy CRs, and potentially to environments where CR scattering is very strongly suppressed since, in both these regimes, the mean free path that CRs travel before becoming isotropised can be large. Fortunately our Itô calculus-based formulation of the problem is readily extensible to include the pitch angle dimension as well, and we intend to include an option to solve problems in the anisotropic regime in future releases.

Third, CRIPTIC does not yet predict neutrino emission. This is a straightforward extension to the existing code, since the main process responsible for producing neutrinos – nuclear inelastic collisions leading to charged pion production – is already included via calculation of secondary electrons and positrons. The limitation is rather that existing neutrino observatories are limited to very high energy neutrinos, where the point raised above applies, i.e. it is unclear that one can assume pitch angle isotropisation for the CRs that drive the observable neutrino emission. Extension to neutrinos is therefore likely to wait until after the extension to follow diffusion in pitch angle is complete.

Finally, we note that CRIPTIC is released under an open source license. Users are encouraged to contribute their own expansions, which can be incorporated into future releases.

ACKNOWLEDGEMENTS

We thank I. Kafexhiu for assistance with cross-section calculations, and C. Federrath for assistance in developing an interface for CRIPTIC to read FLASH simulations. CRIPTIC and its associated support software use the following software packages: ASTROPY¹² (Astropy Collaboration 2013, 2018), NUMPY (Harris et al. 2020), MATPLOTLIB (Hunter 2007), GNU Scientific Library (Galassi et al. 2009), and PCG (O’Neill 2014). This research was undertaken with the assistance of resources and services from the National Computational Infrastructure (NCI), award jh2, which is supported by the Australian Government. This research was funded by the Australian Research Council through its Discovery Projects funding scheme, award DP190101258.

DATA AVAILABILITY

CRIPTIC is available from <https://bitbucket.org/krumholz/criptic/src/master/> under an open source license. The setup files and analysis scripts required to reproduce all the test problems described in this paper are included in the repository. An introduction and users’ guide is available at <https://criptic.readthedocs.io/>.

REFERENCES

- Ambrosone A., Chianese M., Fiorillo D. F. G., Marinelli A., Miele G., 2022, *MNRAS*, 515, 5389
- Armillotta L., Ostriker E. C., Jiang Y.-F., 2021, *ApJ*, 922, 11
- Armillotta L., Ostriker E. C., Jiang Y.-F., 2022, *ApJ*, 929, 170
- Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, 47, 481
- Astropy Collaboration, 2013, *A&A*, 558, A33
- Astropy Collaboration, 2018, *AJ*, 156, 123
- Beattie J. R., Krumholz M. R., Federrath C., Sampson M., Crocker R. M., 2022, preprint ([arXiv:2203.13952](https://arxiv.org/abs/2203.13952))
- Bisbas T. G., Papadopoulos P. P., Viti S., 2015, *ApJ*, 803, 37
- Blumenthal G. R., Gould R. J., 1970, *Rev. Mod. Phys.*, 42, 237
- Chan T. K., Kereš D., Hopkins P. F., Quataert E., Su K. Y., Hayward C. C., Faucher-Giguère C. A., 2019, *MNRAS*, 488, 3716
- Cherenkov Telescope Array Consortium, 2019, *Science with the Cherenkov Telescope Array*. World Scientific Publishing, Singapore
- Crocker R. M., Krumholz M. R., Thompson T. A., 2021a, *MNRAS*, 502, 1312
- Crocker R. M., Krumholz M. R., Thompson T. A., 2021b, *MNRAS*, 503, 2651
- Dirac P. A. M., 1930, *Math. Proc. Camb. Phil. Soc.*, 26, 361
- Draine B. T., 2011, *Physics of the Interstellar and Intergalactic Medium*. Princeton University Press, Princeton, NJ
- Dullemond C. P., 2012, *Astrophysics Source Code Library*, record ascl:1202.015
- Enßlin T. A., Pfrommer C., Springel V., Jubelgas M., 2007, *A&A*, 473, 41
- Evoli C., Gaggero D., Vittino A., Di Bernardo G., Di Mauro M., Ligorini A., Ullio P., Grasso D., 2017, *J. Cosmol. Astropart. Phys.*, 2017, 015
- Fehlberg E., 1970, *Computing*, 6, 61
- Galassi M., Davies J., Theiler J., Gough B., Jungman G., Alken P., Booth M., Rossi F., 2009, *GNU Scientific Library Reference Manual*. 3rd edn
- García-Portugués E., 2022, *Notes for Nonparametric Statistics*. Available at: <https://bookdown.org/egarp/NP-UC3M/>

¹²<http://www.astropy.org>

Gardiner C., 2009, *Stochastic Methods: A Handbook for the Natural and Physical Sciences*, 4 edn. Springer, Berlin

Ginsburg A. et al., 2016, *A&A*, 586, A50

Girichidis P. et al., 2016, *ApJ*, 816, L19

Girichidis P., Pfrommer C., Pakmor R., Springel V., 2022, *MNRAS*, 510, 3917

Gould R. J., 1972, *Physica*, 60, 145

Harris C. R. et al., 2020, *Nature*, 585, 357

Hopkins P. F. et al., 2020, *MNRAS*, 492, 3465

Hopkins P. F., Squire J., Butsky I. S., Ji S., 2021a, preprint ([arXiv:2112.02153](https://arxiv.org/abs/2112.02153))

Hopkins P. F., Chan T. K., Squire J., Quataert E., Ji S., Kereš D., Faucher-Giguère C.-A., 2021b, *MNRAS*, 501, 3663

Hopkins P. F., Butsky I. S., Panopoulou G. V., Ji S., Quataert E., Faucher-Giguère C.-A., Keres D., 2022, *MNRAS*, 516, 3470

Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90

Ivlev A. V., Silsbee K., Padovani M., Galli D., 2021, *ApJ*, 909, 107

Jones F. C., 1968, *Phys. Rev.*, 167, 1159

Kafexhiu E., 2016, *Phys. Rev. C*, 94, 064603

Kafexhiu E., Aharonian F., Taylor A. M., Vila G. S., 2014, *Phys. Rev. D*, 90, 123014

Kelner S. R., Aharonian F. A., Bugayov V. V., 2006, *Phys. Rev. D*, 74, 034018

Kim Y.-K., Rudd M. E., 1994, *Phys. Rev. A*, 50, 3954

Kim Y.-K., Santos J. P., Parente F., 2000, *Phys. Rev. A*, 62, 052710

Kissmann R., 2014, *Astropart. Phys.*, 55, 37

Knudsen H., Brun-Nielsen L., Charlton M., Poulsen M. R., 1990, *J. Phys. B At. Mol. Phys.*, 23, 3955

Kopp A., Büsching I., Strauss R. D., Potgieter M. S., 2012, *Comput. Phys. Commun.*, 183, 530

Krumholz M. R., Crocker R. M., Xu S., Lazarian A., Rosevear M. T., Bedwell-Wilson J., 2020, *MNRAS*, 493, 2817

Lacki B. C., Thompson T. A., 2010, *ApJ*, 717, 196

Lacki B. C., Thompson T. A., Quataert E., 2010, *ApJ*, 717, 1

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

Merten L., Becker Tjus J., Fichtner H., Eichmann B., Sigl G., 2017, *J. Cosmol. Astropart. Phys.*, 2017, 046

Mücke A., Rachen J. P., Engel R., Protheroe R. J., Stanev T., 1999, *PASA*, 16, 160

Narayanan D., Krumholz M. R., 2017, *MNRAS*, 467, 50

Narayanan D. et al., 2021, *ApJS*, 252, 12

O'Neill M. E., 2014, Technical Report HMC-CS-2014-0905, PCG: A Family of Simple Fast Space-Efficient Statistically Good Algorithms for Random Number Generation. Harvey Mudd College, Claremont, CA

Oka T., Geballe T. R., Goto M., Usuda T., Benjamin M. J., Indriolo N., 2019, *ApJ*, 883, 54

Papadopoulos P. P., 2010, *ApJ*, 720, 226

Pattler R. E., 1959, *Q. J. Mech. App. Math.*, 12, 407

Patwary M. M. A. et al., 2016, preprint ([arXiv:1607.08220](https://arxiv.org/abs/1607.08220))

Peretti E., Blasi P., Aharonian F., Morlino G., 2019, *MNRAS*, 487, 168

Peskin M. E., Schroeder D. V., 1995, *An Introduction to Quantum Field Theory*. Addison-Wesley, Boston

Prantzos N. et al., 2011, *Rev. Mod. Phys.*, 83, 1001

Roth M. A., Krumholz M. R., Crocker R. M., Celli S., 2021, *Nature*, 597, 341

Rudd M. E., Kim Y. K., Madison D. H., Gay T. J., 1992, *Rev. Mod. Phys.*, 64, 441

Salem M., Bryan G. L., 2014, *MNRAS*, 437, 3312

Sampson M. L., Beattie J. R., Krumholz M. R., Crocker R. M., Federrath C., Seta A., 2022, preprint ([arXiv:2205.08174](https://arxiv.org/abs/2205.08174))

Schultz D. R., Krstić P. S., Lee T. G., Raymond J. C., 2008, *ApJ*, 678, 950

Skilling J., 1975, *MNRAS*, 172, 557

Socrates A., Davis S. W., Ramirez-Ruiz E., 2008, *ApJ*, 687, 202

Strong A. W., Moskalenko I. V., 1998, *ApJ*, 509, 212

Strong A. W., Moskalenko I. V., Ptuskin V. S., 2007, *Ann. Rev. Nucl. Part. Sci.*, 57, 285

Tanaka K., Nagai M., Kamegai K., 2021, *ApJ*, 915, 79

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

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

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

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

Wibking B. D., Krumholz M. R., 2022, *MNRAS*, in press

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

Yang R.-z., Kafexhiu E., Aharonian F., 2018, *A&A*, 615, A108

Yeast-Hull T. M., Gallagher J. S. I., Zweibel E. G., Everett J. E., 2014, *ApJ*, 780, 137

Yeast-Hull T. M., Gallagher J. S., Zweibel E. G., 2016, *MNRAS*, 457, L29

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

APPENDIX A: DETAILS OF THE TREE ALGORITHM

Here, we describe the algorithm we use to evaluate equation (74) to reconstruct integrals over the CR distribution function and their gradients. For convenience, we rewrite this equation here as

$$q = \sum_{s,i} g(\mathbf{x}_{si}) \Theta_{si} \Upsilon_{si} w_{si}, \quad (\text{A1})$$

where q is one of the quantities – n , P , or U – on the left-hand side of equation (74), w is a corresponding quantity computed from the properties of each CR packet – 1 , $v_{si} p_{si}$, and T_{si} , and $g(\mathbf{x}_{si}) = K_{\mathbf{H}}(\mathbf{x} - \mathbf{x}_{si})$. We shall refer to the quantities q appearing on the left hand as field quantities, and the quantities w appearing on the right-hand side as field weights. Note that the contribution from each packet is therefore the product of a purely geometric term $g(\mathbf{x}_{si})$ that depends only on packet position, a step indicator Θ_{si} that depends only on the relative rigidities of the packet for which the field quantity is being computed and the packets contributing to it, and a term $\Upsilon_{si} w_{si}$ that depends only on other properties of the contributing packet. Our algorithm is based on this decomposition. Further note that equation (76), describing the gradients of field quantities, can be written in a completely analogous fashion, simply by changing the geometric term to $g(\mathbf{x}_{si}) = -\mathbf{H}_{\nabla}^{-1}(\mathbf{x} - \mathbf{x}_{si}) K_{\mathbf{H}_{\Delta}}(\mathbf{x} - \mathbf{x}_{si})$. We can therefore apply the same algorithm to gradients of field quantities, with only very minor modifications that we discuss below.

Given this discussion, we first describe how we construct the kd-tree in Appendix A1, the algorithm we use to evaluate equation (74) in non-distributed memory calculations in Appendix A2, and then the extension to the distributed memory case in Appendix A3.

A1 Building the kd-tree

We construct a balanced kd-tree by standard methods, and assign each leaf a bandwidth tensor $\mathbf{H}$ as described in Section 3.1. We then carry out an additional step: for each leaf L in the tree, we compute $q_{R_n} = \sum_{s,i \in L} \Theta(R_{s,i} - R_n) \Upsilon_{si} w_{si}$ for a series of rigidities R_n uniformly spaced in logarithm from the largest to the smallest packet rigidity present in the volume. In words, this sum represents the total maximum possible contribution that the packets in the leaf could make to a field quantity q for a CR with rigidity R_n .

Once we have evaluated these sums for every leaf, we recursively compute the corresponding sums for every other node in the tree, by simply summing the results from that node's two children. In this way, for every node in the tree we record the maximum possible contribution q_{R_n} that packets contained in that node could make to the field quantities of packets with rigidities $> R_n$.

A2 Shared memory

We evaluate equation (A1) for all packets in the tree by processing one leaf at a time, using a separate OpenMP thread for each leaf. For each CR packet in the leaf, we compute its rigidity R and find the corresponding rigidities R_n that bracket it, i.e. we find n such that $R_n < R < R_{n+1}$. The algorithm we apply relies on the `nodeList`: a list of all the non-leaf nodes of the tree we have examined so far, along with, for each node a central estimate q_{node} and an error e_{node} for the contribution that packets contained in that node make to the field quantities. The central estimates and errors have the property that the true contribution of the packets in each tree node, defined by equation (A1) evaluated for the packets within the node, lies strictly in the range $(q_{\text{node}} - e_{\text{node}}, q_{\text{node}} + e_{\text{node}})$.

(1) Start the algorithm by adding the root node of the tree to `nodeList`. For each CR packet in the leaf whose field quantities are being computed, evaluate the minimum and maximum possible contributions q_{min} and q_{max} to the field quantities made by packets contained within the root node; the minimum is simply given by $q_{\text{min}} = \min(g(\mathbf{x}))q_{R_{n+1}}$, and the maximum by $q_{\text{max}} = \max(g(\mathbf{x}))q_{R_n}$, where here the minimum and maximum of g are evaluated over the bounding box of the node. In words, we find the minimum possible contribution q_{min} by assuming that all the packets in the node are at the location $\mathbf{x}$ that makes the geometric factor $g(\mathbf{x})$ as small as possible, and that the target packet is at the largest possible rigidity, R_{n+1} ; similarly, the maximum possible contribution arises if all the packets in the node are at the location where $g(\mathbf{x})$ has its maximum, and the target packet is at its smallest possible rigidity, R_n . We can then take $q_{\text{node}} = (q_{\text{min}} + q_{\text{max}})/2$ and $e_{\text{node}} = (q_{\text{max}} - q_{\text{min}})/2$. We add these quantities to `nodeList`.

(2) Evaluate the sum of q_{node} and e_{node} over all the nodes in `nodeList` and all packets in the leaf we are processing. Define the maximum possible relative error for each packet in the leaf by $\text{RE} = \sum e_{\text{node}} / (\sum q_{\text{node}} - \sum e_{\text{node}})$, where the maximum is over all packets; the true value of q is guaranteed to differ from $\sum q_{\text{node}}$ by at most a factor of RE for the packet with the largest error. If RE is below a user-specified tolerance for all packets in the leaf, terminate iteration and set $q = \sum q_{\text{node}}$.

(3) If the RE exceeds the tolerance for any packet in the leaf, search through `nodeList` and find the node that makes the largest contribution to RE for the packet with the largest total RE. We then ‘open’ this node by removing it from `nodeList`, and replacing it with its two children. If those children are not leaves, we compute q_{node} and e_{node} for them exactly as for the root node in step 1. If they are leaves, we compute q_{node} for them by evaluating the sum

in equation (A1) directly for the packets those leaves contain, and setting the corresponding error e_{node} to zero.

(4) Go back to step 2, and repeat until RE is below the tolerance for all packets.

When applying this algorithm to the gradients of field quantities, in `nodeList` we have not only estimates of field quantities q_{node} , but estimates of their gradients ∇q_{node} , and the corresponding uncertainties e_{node} and ∇e_{node} . Since ∇q is a vector quantity, we define the relative error as $\text{RE} = |\sum \nabla e_{\text{node}}| h_{\text{max}} / (\sum q_{\text{node}} - \sum e_{\text{node}})$, where h_{max} is the largest eigenvalue of the bandwidth tensor $\mathbf{H}$.

A3 Distributed memory

As described in Section 3.5, in a distributed memory parallel calculation, on each rank we have a tree, some of whose nodes may have children that reside on a different MPI rank. It is therefore possible that, finding the node that makes the largest contribution to the relative errors (step 3 in Appendix A2), we will find that node cannot be opened because its children are only available on another MPI rank. If this occurs, we remove the node from the node list and do not add any children, but we record the MPI rank on which those children live, and we keep track of the total central estimate q_{ext} and uncertainty e_{ext} contributed by such ‘external’ nodes, and we define the relative error to include their contributions: $\text{RE} = (\sum e_{\text{node}} + e_{\text{ext}}) / (\sum q_{\text{node}} - \sum e_{\text{node}} - e_{\text{ext}})$, and similarly for gradients. We distinguish this from the ‘local’ relative error $\text{RE}_{\text{loc}} = \sum e_{\text{node}} / (\sum q_{\text{node}} - \sum e_{\text{node}})$.

If at any point while iterating, we reach a state where RE_{loc} is below our target tolerance for all packets, but RE exceeds our tolerance for at least some of them, we conclude that we cannot complete those packets without access to the information contained on other MPI ranks. We therefore send an MPI request to all the ranks that contribute to e_{ext} , and request that they evaluate their contribution to the target packets; they do so using the same algorithm as given in Appendix A2, but rather than starting from the root node of the tree, they start from the roots of their own local trees, i.e. evaluating only their own contributions. Once they have computed these contributions with sufficient accuracy that the total estimate will be below the tolerance, they send these estimates back to the rank that made the request. This rank then sums the contributions from all external ranks to arrive at a final estimate for the field quantities.

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